arXiv:1601.03945v1 [cs.CV] 15 Jan 2016

Improved graph-based SFA: Information preservation complements the slowness principle

Alberto N. Escalante-B.*, Laurenz Wiskott

Institut für Neuroinformatik
Ruhr-University Bochum, Germany

Abstract

Slow feature analysis (SFA) is an unsupervised-learning algorithm that extracts slowly varying features from a multi-dimensional time series. A supervised extension to SFA for classification and regression is graph-based SFA (GSFA). GSFA is based on the preservation of similarities, which are specified by a graph structure derived from the labels. It has been shown that hierarchical GSFA (HGSFA) allows learning from images and other high-dimensional data. The feature space spanned by HGSFA is complex due to the composition of the nonlinearities of the nodes in the network. However, we show that the network discards useful information prematurely before it reaches higher nodes, resulting in suboptimal global slowness and an under-exploited feature space.

To counteract these problems, we propose an extension called hierarchical information-preserving GSFA (HiGSFA), where information preservation complements the slowness-maximization goal. We build a 10-layer HiGSFA network to estimate human age from facial photographs of the MORPH-II database, achieving a mean absolute error of 3.50 years, improving the state-of-the-art performance. HiGSFA and HGSFA support multiple-labels and offer a rich feature space, feed-forward training, and linear complexity in the number of samples and dimensions. Furthermore, HiGSFA outperforms HGSFA in terms of feature slowness, estimation accuracy and input reconstruction, giving rise to a promising hierarchical supervised-learning approach.

Keywords: Supervised dimensionality reduction, Similarity-based learning, Information preservation, Deep networks, Age estimation

1. Introduction

Supervised dimensionality reduction (supervised DR) has promising applications in computer vision, pattern recognition and machine learning, where

*Corresponding author.
Email addresses: alberto.escalante@ini.rub.de (Alberto N. Escalante-B.), laurenz.wiskott@ini.rub.de (Laurenz Wiskott)

it is frequently used as a pre-processing step to solve classification and regression problems with high-dimensional input. The goal of supervised DR is to extract a low-dimensional representation of the input samples that contains the predictive information about the labels (e.g., [1]). One advantage is that dimensions irrelevant to the supervised problem can be discarded, resulting in a more compact representation and more accurate estimations.

The slowness principle requires the extraction of the most slowly changing features, see [2], and can be used for DR by itself or before other algorithms. It has been shown that this principle might explain in part how the neurons in the brain self-organize to compute invariant representations. Slow features can be computed using a few methods, such as online learning rules (e.g., [3, 4]), slow feature analysis (SFA) [5, 6], which is a closed-form algorithm specific for this task that has biologically feasible variants [7], and an incremental-learning version (inc-SFA) [8].

The optimization objective of SFA is the minimization of the squared output differences between (temporally) consecutive pairs of samples. Although SFA is unsupervised, the order of the samples can be regarded as a weak form of supervised information, allowing the use of SFA for supervised learning tasks.

Graph-based SFA (GSFA) [9, 10] is an extension of SFA explicitly designed for the solution of classification and regression problems and yields more accurate label estimations than SFA. GSFA has been used to estimate age and gender from synthetic face images [11], and more recently for traffic sign classification [12] and face detection [13]. GSFA is trained with a training graph, where the vertices are the samples and the edges represent similarities of the corresponding labels. The optimization objective of GSFA is the minimization of weighted squared output differences between samples connected by an edge, where the edges are carefully chosen to reflect label similarities.

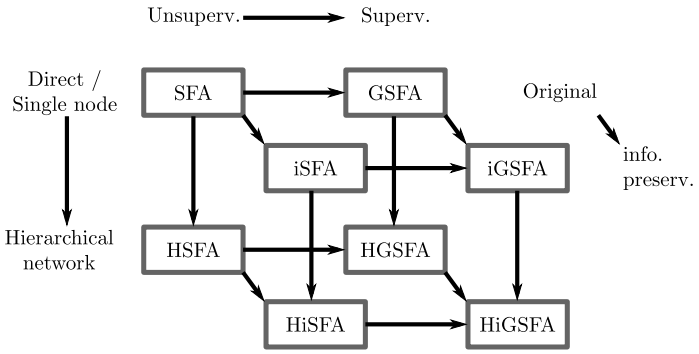

Figure 1: The combination of three extensions of SFA (graph-based, hierarchical, information-preserving) gives rise to 8 different versions of SFA. In this article we propose the versions with information preservation, being HiGSFA the most significant one.

A common problem in machine learning is the high computational cost of algorithms when the data are high-dimensional. This is also the case when GSFA

is applied to the data directly (direct GSFA), because it has cubic complexity w.r.t. the number of dimensions. However, processing high-dimensional data is still practical if one resorts to hierarchical GSFA (HGSFA). Instead of extracting slow features in a single step, the features are computed by applying GSFA on lower dimensional data chunks repeatedly. Other advantages of HGSFA over direct GSFA include lower memory requirements and a complex global nonlinearity due to the composition of nonlinear transformations through the layers. Moreover, the local extraction of slow features from lower-dimensional data chunks typically results in less overfitting.

In this article, we show that HGSFA suffers from a drawback: The separate instances of GSFA that process the low-dimensional data chunks in the network, also called GSFA nodes, may prematurely discard features that are not slow at a local level but that would have been useful to improve global slowness (i.e., the slowness of the final output features) if combined with features from other nodes. This drawback, which we call unnecessary information loss, leads to a suboptimal use of the feature space comprised by the network and also affects the estimation accuracy for the underlying supervised learning problem.

Principle or heuristic	**Implemented through**
Slowness principle, exploit label similarities through slowness	GSFA, supervised training graphs
Spatial localization of features, divide and conquer approach	Hierarchical processing
Robustness to outliers	Normalized or saturating nonlinear expansions
Information preservation (new)	Minimization of a reconstruction error, PCA
Multiple information channels (new)	Multi-label learning combining efficient training graphs

Table 1: Principles, heuristics and ideas considered in HiGSFA and the base methods or algorithms used to exploit them.

To reduce the unnecessary information loss in HGSFA, we propose to complement slowness with information preservation (i.e., maximization of mutual information between the input data and the output features). For simplicity and efficiency, we implement this idea as the minimization of a reconstruction error. The resulting network is called hierarchical information-preserving GSFA (HiGSFA), and the algorithm constituting each node of the network is called information-preserving GSFA (iGSFA). The features computed by iGSFA can be divided in two parts: A slow part, which is a linear transformation of the (nonlinear) features computed with GSFA, and an input-reconstructive part. To compute the input-reconstructive part the input data are approximated using the slow part, resulting in residual data that are then processed by PCA. The construction ensures that both parts are decorrelated and the features have com-

patible scales. Different versions of SFA with and without information preservation are shown in Figure 1. The principles and heuristics behind HiGSFA are presented in Table 1.

The experiments show the advantages of HiGSFA over HGSFA: (1) slower features, (2) better generalization to unseen data, (3) much better input reconstruction (see Figure 2), and (4) improved accuracy for the supervised learning problem. Furthermore, the computational and memory requirements of HiGSFA are asymptotically the same as in HGSFA.

Figure 2: (a) An image from a private database after pose normalization. (b) The same image fully pre-processed (i.e., after pose normalization and face sampling, 96×96 pixels). Linear reconstructions on 75 features extracted with either (c) PCA, (d) HiGSFA or (e) HGSFA. (f) Average over all pre-processed images of the MORPH-II database.

2. Related work

HiGSFA is the main extension to SFA proposed in this work and belongs to supervised DR. Other algorithms for supervised DR include Fisher discriminant analysis (FDA) [14], local FDA (LFDA) [15], pairwise constraints-guided feature projection (PCGFP) [16], semi-supervised dimensionality reduction (SSDR) [17], and semi-supervised LFDA (SELF) [18].

Existing extensions to SFA include extended SFA (xSFA) [19], generalized SFA (genSFA) [20] and graph-based SFA (GSFA) [9, 10, 11]. HiGSFA extends hierarchical GSFA (HGSFA) by adding information preservation. With some adaptations, SFA has been shown to be successful for classification (e.g., [21, 22, 23, 24, 25]), and regression (e.g., [21]).

From a mathematical point of view, SFA, LPP, genSFA and GSFA belong to the same family of optimization problems, see [26], and can be solved via generalized eigenvalues. Conversions between these four algorithms are possible. Two differences between LPP and GSFA are that in GSFA the vertex weights are independent of the edge weights and that GSFA is invariant to the scale of the weights, providing a normalized objective function.

In spite of being closely related mathematically, SFA, LPP, genSFA and GSFA originate from different backgrounds and were first motivated with different goals in mind. Therefore, they usually have different applications and are trained with different similarity matrices, resulting in features with different properties. For instance, LPP originates from the field of manifold learning and transitions are typically defined from input similarities (nearest neighbors). SFA originates from unsupervised learning and transitions are typically defined by the temporal ordering of the samples. genSFA typically combines input similarities with class information. In GSFA, transitions typically reflect label similarities. For SFA and GSFA the use of input similarities might be disadvantageous, because it might compromise the invariance of the features, which is one of their central goals.

HiGSFA is a hierarchical implementation of information-preserving GSFA (iGSFA), which has one parameter Δ_T (Δ-threshold) more than GSFA. This parameter balances the number of slow and reconstructive features. When $\Delta_T < 0$ iGSFA becomes equivalent to PCA, and when $\Delta_T \geq 4.0$ iGSFA computes a linear transformation of GSFA features. Theory justifies fixing Δ_T slightly smaller than 2.0 (Section 4.4).

The rest of the article is organized as follows. In the next section we review GSFA. In Section 4, we analyze the advantages and limitations of hierarchical networks for slow feature extraction. In Section 5, we propose the iSFA (and iGSFA) algorithm. In Section 6, we evaluate HiGSFA experimentally using the problem of age estimation from facial photographs. We conclude with a discussion section.

3. Overview of Graph-based SFA (GSFA)

A training graph $G = (\mathbf{V}, \mathbf{E})$ (illustrated in Figure 3.a) has a set $\mathbf{V}$ of vertices $\mathbf{x}(n)$, each vertex being a sample (i.e. an I-dimensional vector), and a set $\mathbf{E}$ of edges $(\mathbf{x}(n), \mathbf{x}(n'))$, which are pairs of samples, with $1 \leq n, n' \leq N$. The index n (or n') substitutes the time variable t of SFA. The edges are undirected and have symmetric weights $\gamma_{n,n'} = \gamma_{n',n}$, which indicate the similarity between the connected vertices; also each vertex $\mathbf{x}(n)$ has an associated weight $v_n > 0$, which can be used to reflect its frequency. This representation includes the standard time series of SFA as a special case (Figure 3.b).

The GSFA optimization problem [10] can be stated as follows. For $1 \leq j \leq J$, find features $y_j(n) = g_j(\mathbf{x}(n))$ with $g_j \in \mathcal{F}$, where $1 \leq n \leq N$ and $\mathcal{F}$ is the feature space, such that the objective function (weighted delta value)

$$\Delta_j \stackrel{\text{def}}{=} \frac{1}{R} \sum_{n,n'} \gamma_{n,n'} (y_j(n') - y_j(n))^2 \text{ is minimal} \tag{1}$$

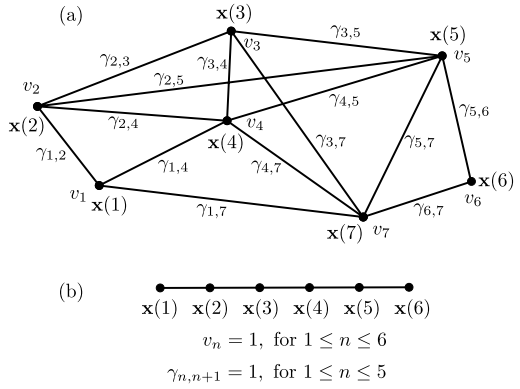

Figure 3: (a) Example of a training graph with $N = 7$ vertices. (b) A linear graph suitable for GSFA that learns the same features as SFA on the time series $\mathbf{x}(1), \ldots, \mathbf{x}(6)$. (Figure from [10]).

under the constraints

$$\frac{1}{Q} \sum_n v_n y_j(n) = 0\,, \tag{2}$$

$$\frac{1}{Q} \sum_n v_n (y_j(n))^2 = 1\,, \text{ and} \tag{3}$$

$$\frac{1}{Q} \sum_n v_n y_j(n) y_{j'}(n) = 0 \text{ for } j' < j \tag{4}$$

with $R \stackrel{\text{def}}{=} \sum_{n,n'} \gamma_{n,n'}$ and $Q \stackrel{\text{def}}{=} \sum_n v_n$.

The constraints (2–4) are called weighted zero mean, weighted unit variance, and weighted decorrelation. The factors $1/R$ and $1/Q$ provide invariance to the scale of the edge weights as well as to the scale of the vertex weights. Typically, a linear feature space is used, but the input samples are preprocessed by a nonlinear expansion function.

Depending on the training graph chosen, the properties of the features extracted by GSFA can be quite different. Training graphs for classification favor connections between samples from the same class, whereas graphs for regression favor connections between samples with similar labels. In Section 6.3, we show how to combine graphs for classification and regression into an efficient graph to learn age (regression), gender (classification) and race (classification) simultaneously.

The motivation for HiSFA and HiGSFA is to correct certain disadvantages of HSFA and HGSFA while preserving their advantages. An analysis of these advantages and disadvantages is the topic of the next section.

4. Advantages and limitations of HSFA networks

In this section we analyze HSFA networks from the point of view of their advantages, particularly their computational complexity, and their limitations, particularly unnecessary information loss. We focus on HSFA but HGSFA is covered by extension. Understanding the drawbacks of HSFA is crucial to justify the extensions with information preservation proposed in Section 5.

4.1. Previous work on HSFA and terminology

The first example of HSFA appears also in the paper that first introduces the SFA algorithm [5], where it is employed as a model of the visual system for learning invariant representations.

Various contributions have continued with this biological interpretation. In [27], HSFA has been used to learn invariant features from the simulated view of a rat that moves inside a box. In conjunction with a sparseness post-processing step, the extracted features are similar to the responses of place cells in the hippocampus.

Other works have focused more on its computational efficiency compared to direct SFA. In [28], HSFA has been used for object recognition from images and to estimate pose parameters of single objects moving and rotating over a homogeneous background. In [10], an HGSFA network with 11 layers has been used to accurately find the horizontal position of faces in photographs, which is a sub-problem of face detection.

In general, HSFA networks should be directed and acyclic with arbitrary edges otherwise, but usually they are composed of multiple layers, each having a regular structure. Most of the HSFA networks in the literature have a similar composition. Typical differences are how the data are split into smaller chunks and the particular processing done by the SFA nodes themselves.

The structure of the network usually follows the structure of the data. For instance, networks for data in 1 dimension (e.g., audio data represented as fixed length vectors) typically have a 1-dimensional structure (e.g., Figure 4), and networks for data in 2 dimensions (e.g., images) have a 2-dimensional structure. This idea extends to voxel data in 3 dimensions, and beyond.

For simplicity, we refer to the input data as layer 0. Important parameters that define the structure of a network include: (a) The *output dimensionality* of the nodes. (b) The *fan-in* of the nodes, which is the number of nodes (or data elements) in a previous layer that feed them. (c) The *receptive fields* of the nodes, which refer to all the elements of the input data that (directly or indirectly) provide input to a particular node. (d) The *stride* of a layer, which tells how far apart the inputs to adjacent nodes in a layer are. If the stride is smaller than the fan-in, then at least one node in the previous layer will feed two or more nodes in the current layer. This is called *receptive field overlap*.

4.2. Advantages of HSFA

Interestingly, hierarchical processing reduces overfitting and can be seen as a regularization method. Generalization is improved even when a larger number

of parameters have to be learned. This can be explained by the fact that the input to each node in the hierarchical network has the same number of samples as the original input but much smaller dimensionality, frequently leading to good generalization to unseen data in individual nodes as well as in the whole network. The difference in generalization between HSFA and direct SFA is more evident when polynomial expansions are involved, because the larger expanded dimensionality in direct SFA translates into stronger overfitting.

Another interesting property of HSFA is that the nonlinearity of the nodes accumulates across the layers, so that even when using simple expansions the network as a whole may describe a highly nonlinear feature space [29]. Depending on the data, such a complex feature space may be necessary to extract the slowest hidden parameters and to solve the supervised problem with good accuracy.

A key motivation for preferring HSFA over direct SFA is its computational efficiency. We focus here on the complexity of training rather than of feature extraction, because the former is more relevant for applications, and the latter is relatively lightweight in both cases. Training linear SFA has a computational (time) complexity

$$T_{\mathrm{SFA}}(N, I) = \mathcal{O}(NI^2 + I^3)\,, \tag{5}$$

where N is the number of samples and I is the input dimensionality (possibly after a nonlinear expansion). The same complexity holds for GSFA if one uses an efficient training graph (e.g., the clustered or serial graphs, see Section 6.3), otherwise it can be as large as (for arbitrary graphs)

$$T_{\mathrm{GSFA}}(N, I) = \mathcal{O}(N^2I^2 + I^3)\,. \tag{6}$$

For large I (i.e., high-dimensional data) direct SFA and GSFA are therefore inefficient[1]. Their complexity can be reduced by using HSFA and HGSFA. The exact complexity depends on the structure and parameters of the hierarchical network. As we prove below, it can be linear in I and N for certain networks.

4.3. Complexity of a quadratic HSFA Network

Although existing applications have exploited the speed of HSFA, we are not aware of any analysis of its actual complexity. We compute the computational complexity of a concrete quadratic HSFA (QHSFA) network for 1D data with L layers, which is illustrated in Figure 4. Each node of the network performs quadratic SFA (QSFA, i.e., a quadratic expansion followed by linear SFA). The receptive field, fan-in, and stride of the nodes in layer 1 is k input values, where k is a fixed parameter. In the rest of the layers the fan-in and stride of the nodes is 2 nodes. Assume that the total input dimensionality is I and that every node reduces the dimensionality of its input from k components to $k/2$ components.

[1]The problem is still feasible if N is small enough so that one might apply singular value decomposition methods. However, a small number of samples $N < I$ usually results in pronounced overfitting.

From the structure described above, it follows that the receptive fields of the nodes are non-overlapping, the network's output has $k/2$ components, the number of layers L is related to I and k:

$$I \;=\; k2^{L-1}\,, \tag{7}$$

and the total number of nodes in the network is

$$M \;=\; 2^L - 1\,. \tag{8}$$

The input to each QSFA node has k dimensions, which is increased to $k(k+3)/2$ dimensions by the quadratic expansion. Afterwards, linear SFA reduces the dimensionality to $k/2$. Hence, the complexity of training a single node is

$$T_{\mathrm{QSFA}}(N,k) \overset{(5)}{=} \mathcal{O}(N(k(k+3)/2)^2 + (k(k+3)/2)^3) \tag{9}$$

$$= \mathcal{O}(Nk^4 + k^6)\,. \tag{10}$$

The number of nodes is $M \overset{(7,8)}{=} \mathcal{O}(I/k)$. Therefore, the complexity of training the whole network is

$$T_{\mathrm{QHSFA}}(N,I,k) \overset{(10)}{=} \mathcal{O}((Nk^4 + k^6)I/k) = \mathcal{O}(INk^3 + Ik^5)\,. \tag{11}$$

This means the complexity of the QHSFA network above is linear w.r.t. the input dimension I, whereas the complexity of direct QSFA is

$$T_{\mathrm{QSFA}}(N,I) \overset{(10)}{=} \mathcal{O}(NI^4 + I^6)\,, \tag{12}$$

which is linear w.r.t. I^6. Thus, the QHSFA network has a large computational advantage over direct QSFA.

Since each layer in the QHSFA network is quadratic, in general the output features of layer l can be written as polynomials of degree 2^l. In particular, the output features of the network are polynomials of degree 2^L. However, the actual feature space does not span all the polynomials of this degree but only a subset of them due to the restricted connectivity of the network. In contrast, direct QSFA only contains quadratic polynomials (although all of them).

One could train direct SFA on data expanded by a polynomial expansion of degree 2^L. The expanded dimensionality would be $\sum_{d=0}^{2^L} \binom{d+I-1}{I-1} = \mathcal{O}(I^{2^L})$, resulting in a complexity of $\mathcal{O}(NI^{2^{L+1}} + I^{3\cdot 2^L})$, being even less feasible than direct QSFA.

The space (memory) complexity of linear SFA is

$$S_{\mathrm{SFA}}(N,I) = \mathcal{O}(I^2 + NI) \tag{13}$$

(where the term NI is due to the input data). One can reduce this by using HSFA and training each node separately, one at a time. For instance, the space complexity of direct QSFA is

$$S_{\mathrm{QSFA}}(N,I) = \mathcal{O}(I^4 + NI)\,, \tag{14}$$

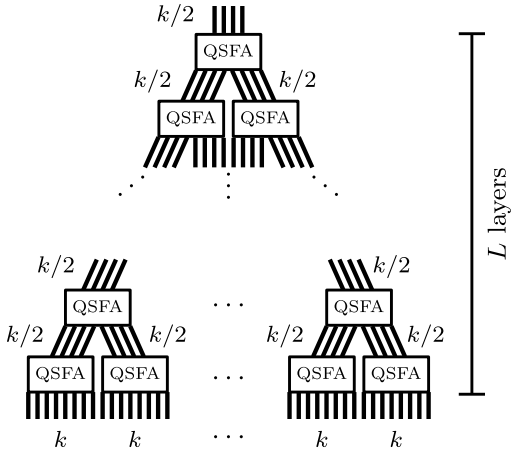

Figure 4: Example of an 1D QHSFA network with binary fan-ins and no overlap. Each node performs quadratic SFA reducing the dimensionality from k to $k/2$ components. A small fan-in results in a network with a large number L of layers, which may be useful to build deep networks.

whereas the space complexity of the QHSFA network is only

$$S_{\text{QHSFA}}(N, I, k) = \mathcal{O}(k^4 + NI)\,. \tag{15}$$

It is possible to design more sophisticated networks than the QHSFA one and preserve a similar computational complexity. For example, the network proposed in Section 6 has overlapping receptive fields, larger fan-ins, a 2D structure, and a complexity also linear in I and N.

4.4. Limitations of HSFA networks

In this section, we analyze the limitations of HSFA networks, or in general, any network in which the nodes have only one criterion for DR, namely, slowness maximization. Otherwise the nodes are unrestricted; they might be linear or nonlinear, include additive noise, clipping, various passes of SFA, etc.

In spite of the remarkable advantages of HSFA, we show that relying only on slowness to determine which aspects of the data that are preserved results in three disadvantages: unnecessary information loss, poor input reconstruction and feature garbling.

Unnecessary information loss. This disadvantage occurs when the nodes discard dimensions of the data that are not significantly slow locally (i.e., at the node level), but which would have been useful for slowness optimization at higher layers of the network if they had been preserved.

We show through a toy experiment that dimensions important for global slowness are not necessarily slow locally. Consider four zero-mean, unit-variance signals: $s_1(t)$, $s_2(t)$, $s_3(t)$ and $n(t)$ that can only take binary values in $\{-1, +1\}$

(n stands for noise here), where $\Delta_{s_1} < \Delta_{s_2} < \Delta_{s_3} < \Delta_n = 2.0$, and all features are independent of the rest[2]. Assume the 4-dimensional input is $(x_1, x_2, x_3, x_4) \stackrel{\text{def}}{=} (s_2, s_1 n, s_3, n)$ and the number of samples is large enough. Thus, the slowest possible feature described by the data is $x_2 x_4 = (s_1 n)n = s_1$ (or equivalently $-x_2 x_4$).

QSFA would be able to extract the slowest feature from these data, but let us assume that a 2-layer QHSFA network with 3 nodes is used, where the output feature is: $\text{QSFA}\big(\text{QSFA}(s_2, s_1 n), \text{QSFA}(s_3, n)\big)$. Each QSFA node reduces the number of dimensions from 2 to 1. Since $\Delta_{s_2} < \Delta_{s_1 n} = 2.0$, the first bottom node computes $\text{QSFA}(s_2, s_1 n) = s_2$, and since $\Delta_{s_3} < \Delta_n = 2.0$, the second bottom node computes $\text{QSFA}(s_3, n) = s_3$. The top node extracts $\text{QSFA}(s_2, s_3) = s_2$. Therefore, the network misses the slowest feature, s_1, even though it belongs to the feature space spanned by the network.

The problem can be expressed in information theoretic terms:

$$I(s_1 n, s_1) = 0\,, \text{ and} \tag{16}$$

$$I(n, s_1) = 0\,, \text{ but} \tag{17}$$

$$I((s_1 n, n), s_1) = H(s_1) > 0\,, \tag{18}$$

where I denotes mutual information, and H is entropy. Thus, it is impossible to locally rule out that a feature contains information that might yield a slow feature, unless one (globally) observes the whole data available to the network.

Unnecessary information loss can also affect applications in practice. For example, Figure 5 shows the Δ values of the slowest features extracted by the first layer of an HGSFA network trained for age estimation on human face images. Most Δ values are approximately 2.0, and only a few of them are less than 2.0.

A feature with $\Delta = 2.0$ can be a transformation of the input data, a transformation of inherent noise, or a mixture of both. In fact, if two or more feasible features have the same Δ value, SFA outputs an arbitrary rotation of them, even though one might prefer features that are transformations of the input only rather than noise. Due to the DR only a few features with $\Delta = 2.0$ may be preserved. The discarded features might still contain useful information, giving place to information loss.

Poor input reconstruction. The goal of input reconstruction is to generate an input from a given slow feature representation. This interesting task has been studied in [30], and may help to determine which features the network is sensitive to, or to find inputs with certain properties.

In the field of image processing, reconstruction might be relevant for image morphing and interpolation. Here, we consider morphing as the task of finding how the input must be modified to reflect modifications introduced to the output features. For example, assume SFA was trained to extract age from facial

[2]One can show that the expected Δ value of a random unit-variance i.i.d. noise feature is 2.0 [12]. The same holds for GSFA if the graph is consistent and has no self-loops.

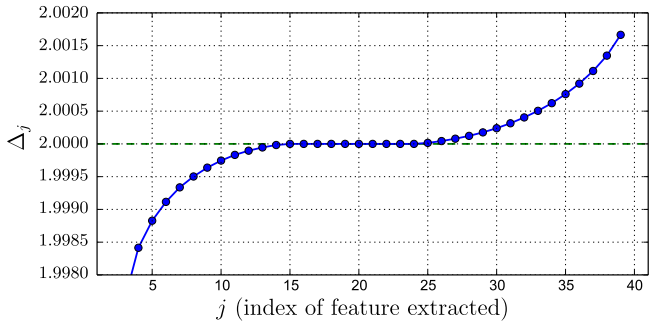

Figure 5: Δ values of the first 40 slow features averaged over all the nodes of the first layer of an HGSFA network trained for age estimation ($\Delta_1 = 1.859$, $\Delta_2 = 1.981$, and $\Delta_3 = 1.995$, not shown). The training graph employed is a serial graph (see Section 6.3) with 32 groups. Most Δ values are close to 2.0, indicating that at this early stage, where the nodes have small (6×6-pixel) receptive fields, the slowest features found are not significantly slow.

images. Morphing would allow us to visualize how a particular person would look a few years older or younger.

Experiments have shown that input reconstruction from the top-level features extracted by HSFA is a challenging task [30]. We have also conducted experiments with various nonlinear methods for input reconstruction, including local and global methods, confirming its difficulty.

We show here that the cause of poor input reconstruction may not be the weakness of the reconstruction algorithms employed but insufficient reconstructive information in the slow features. The extracted features ideally depend only on the hidden slow parameters and are invariant to any other factor. In the age estimation example, if the extracted features are close to the features predicted from theory, they would be strongly related to the age and harmonic functions of it. Thus, they would be mostly invariant to other factors, such as the identity of the person, his facial expression, the background, etc. Therefore, in theory only the age would be available for reconstruction.

In practice, the features usually contain residual information about the input data. However, one cannot rely on this information because it may be partial, making reconstructions not unique, and highly nonlinear, making it difficult to untangle it. Even the features extracted with linear SFA typically result in inaccurate reconstructions. Using many layers of SFA may make the problem more serious in HSFA.

One exception where reconstruction is possible is when SFA is trained with artificial data generated only from slowly changing parameters. In this case, the slow output features should encode the generative parameters, allowing input reconstruction if the reconstruction method is powerful enough.

Feature garbling. Even if a node does not incur in information loss, the features might still not be represented in a useful way for the extraction of slow features in the next nodes. Feature garbling occurs when the features extracted by SFA are

more complex (e.g., highly nonlinear) than necessary without being slower. This complicates the extraction of slow features in higher layers, because the feature space available to the nodes might not suffice to transform the representation into slow features.

We exemplify feature garbling with another toy experiment. Consider two feature vectors $\mathbf{x}(t)$ and $\mathbf{x}'(t) = E_k(\mathbf{x}'(t))$, where E_k is a (bijective) encryption function for some unknown key k. Assume that $\Delta(x_1(t)) = \cdots = \Delta(x_I(t)) = 2.0$, $\Delta(x'_1(t)) = \cdots = \Delta(x'_I(t)) = 2.0$, and that $\mathbf{x}(t)$ and $\mathbf{x}'(t)$ are valid output signals that belong to the feature space. Since both signals have the same Δ-values, SFA might output any of them. The signal $\mathbf{x}(t)$ might be more useful for slowness maximization in the next layers because its features are simpler and more directly connected to the inputs, where as $\mathbf{x}'(t)$ might be useless for the next layers, unless the decryption function belongs to the feature space.

Feature garbling results in information loss. An informative but garbled feature is likely to be discarded later in the hierarchy. One might attempt to reduce feature garbling by using a mild nonlinear expansion in SFA. However, this might compromise the slowness of the extracted features. Similarly, one might try to preserve a large number of features to reduce information loss. However, this might be impractical because it would increase the computational cost and contradicts the goal of dimensionality reduction.

The problems of poor input reconstruction and unnecessary information loss are also connected. This is evident if one distinguishes between two types of information: (a) information about the full input data and (b) information about the global slow parameters. Losing (a) results in poor input reconstruction, whereas losing (b) results in unnecessary information loss. Of course, (a) contains (b). Therefore, both problems originate from losing different but related types of information.

Feature garbling is a theoretical open issue that still needs to be formalized and whose relevance in practice is unknown. Therefore, we do not attempt to counteract it with the proposed extensions.

5. Information-preserving SFA (iSFA)

In this section, we propose information-preserving SFA (iSFA) to counteract the problems of unnecessary information loss and poor input reconstruction. We write iSFA with lowercase 'i' to distinguish it from independent SFA (ISFA) [31]. The extension can also be applied to GSFA and is then called information-preserving GSFA (iGSFA). For simplicity, we first focus on iSFA.

iSFA combines two learning principles: the slowness principle and information preservation without compromising the former in any way. Information preservation requires the maximization of the mutual information between the output features and the input data. However, for finite, discrete and typically unique data samples it is difficult to measure and maximize mutual information unless one assumes some probability model. Therefore, we interpret information preservation more practically and minimize a reconstruction error. A closely related concept is the explained variance, but we avoid this term because it is

typically restricted to linear transformations. When iSFA is used in hierarchical networks, the concept of information preservation might also be called information propagation.

In the rest of the section, first we present a high-level description of the algorithm, then, we describe its construction in detail, show how to approximate an inverse transformation, and discuss its main properties.

5.1. Algorithm overview

The goal of iSFA is to improve the feature extraction of HSFA networks by improving feature extraction at the node level, which is achieved by replacing the SFA nodes with iSFA nodes. Therefore, the structure of the HSFA network can be preserved.

An essential property of iSFA is that the feature vectors are composed of two parts: (1) a slow part derived from SFA features, and (2) a reconstructive part derived from principal components (PCs).

Roughly speaking, the slow part captures the slow aspects of the data and is basically composed of standard SFA features, except for an additional linear mixing step explained in Sections 5.2 and 5.4. The reconstructive part ignores the slowness criteria and instead focuses on describing the input linearly as closely as possible (disregarding the part already described by the slow part). In Section 6 we show that, although the reconstructive features are not particularly slow, they contribute to global slowness maximization.

The proposed algorithm takes care of the following important considerations. (a) Given the output dimension D, it decides how many features the slow and reconstructive part should contain. (b) It minimizes the redundancy between the slow and the reconstructive part, allowing the output features to be more compact and have higher information content. (c) It corrects the amplitudes of the slow features (SFA features usually have unit variance) to make them compatible with the PCs (PCA is a rotation and projection, preserving thus the amplitude of the original data).

5.2. Algorithm description (training phase)

In this section, we present iSFA in detail, or more precisely, its training phase. Figure 6 shows the components involved in the algorithm, and Algorithm 1 gives a concise description.

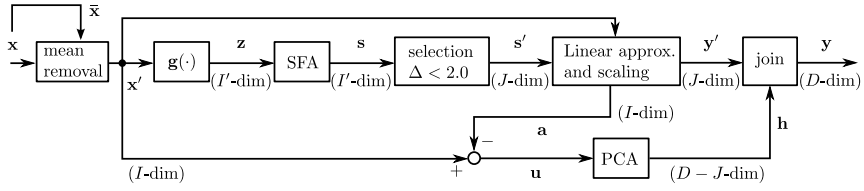

Figure 6: Block diagram of the iSFA node showing the components used for training, feature extraction, and linear input reconstruction. The blocks and signals are explained in the text.

Let $\mathbf{X} \stackrel{\text{def}}{=} (\mathbf{x}_1, \ldots, \mathbf{x}_N)$ be the I-dimensional training data, D the output dimensionality, $g(\cdot)$ the expansion function, and $\Delta_T = 2.0$ (a Δ-threshold, in practice slightly smaller than 2.0). First, the average sample $\bar{\mathbf{x}} \stackrel{\text{def}}{=} \frac{1}{N}\sum_n \mathbf{x}_n$ is removed from the N training samples resulting in the centered data $\mathbf{X}' \stackrel{\text{def}}{=} \{\mathbf{x}'_n\}$, with $\mathbf{x}'_n \stackrel{\text{def}}{=} \mathbf{x}_n - \bar{\mathbf{x}}$, where $1 \le n \le N$. Then, $\mathbf{X}'$ is expanded via $\mathbf{g}(\cdot)$, resulting in vectors $\mathbf{z}_n = \mathbf{g}(\mathbf{x}'_n)$ of dimensionality I'. Afterwards, a new instance of linear SFA is created and trained with the expanded data $\mathbf{Z} \stackrel{\text{def}}{=} \{\mathbf{z}_n\}$. The slow features extracted from the expanded data are denoted $\mathbf{s}_n$. The first J components of $\mathbf{s}_n$ with $\Delta < \Delta_T$ and $J < J' \stackrel{\text{def}}{=} \min(I', D)$, are denoted $\mathbf{s}'_n$. The remaining slow features are discarded.

The next steps correct the amplitude of the J slow features in $\mathbf{S}' \stackrel{\text{def}}{=} \{\mathbf{s}'_n\}$, which have unit variance. The centered data $\mathbf{X}'$ is approximated linearly from $\mathbf{S}'$ by computing a matrix $\mathbf{M}$ and a vector $\mathbf{b}$, as follows.

$$\mathbf{A} \stackrel{\text{def}}{=} \mathbf{M}\mathbf{S}' + \mathbf{b}\mathbf{1}^T \approx \mathbf{X}' \,, \tag{19}$$

where $\mathbf{A}$ is the contribution of the slow part to the approximation of the centered data (i.e., $\mathbf{x}'_n \approx \mathbf{a}_n \stackrel{\text{def}}{=} \mathbf{M}\mathbf{s}'_n + \mathbf{b}$) and $\mathbf{1}$ is a vector of 1s of length N. Since $\mathbf{X}'$ and $\mathbf{S}'$ are centered, $\mathbf{b}$ could be discarded because $\mathbf{b} = \mathbf{0}$. However, when GSFA is used the slow features have *weighted* zero mean, and $\mathbf{b}$ might improve the approximation of $\mathbf{X}'$. The QR decomposition of $\mathbf{M}$ is computed as

$$\mathbf{M} = \mathbf{Q}\mathbf{R} \,, \tag{20}$$

where $\mathbf{Q}$ is orthonormal and $\mathbf{R}$ is upper triangular. Then, the slow feature part is computed as

$$\mathbf{y}'_n = \mathbf{R}\mathbf{s}'_n \,. \tag{21}$$

Section 5.4 justifies this mixing and scaling of the slow features $\mathbf{s}'_n$.

To obtain the reconstructive part, the residual data $\mathbf{u}_n \stackrel{\text{def}}{=} \mathbf{x}'_n - \mathbf{a}_n$ are computed, i.e., the data that remains after the data linearly reconstructed from $\mathbf{y}'_n$ (or $\mathbf{s}'_n$) is removed from the centered data. Afterwards, a new PCA instance is trained with $\{\mathbf{u}_n\}$. Then, the reconstructive part $\mathbf{h}_n$ is defined as the $D - J$ first principal components of $\mathbf{u}_n$.

Afterwards, the concatenation of the slow part $\mathbf{y}'_\mathbf{n}$ (J features) and reconstructive part $\mathbf{h_n}$ ($D - J$ features) results in the D-dimensional output features $\mathbf{y_n} \stackrel{\text{def}}{=} \mathbf{y}'_\mathbf{n}|\mathbf{h_n}$, where $|$ is vector concatenation.

Finally, the algorithm returns $\mathbf{Y} = (\mathbf{y_1}, \ldots, \mathbf{y_N})$, $\bar{\mathbf{x}}$, J, $\mathbf{Q}$, $\mathbf{R}$, $\mathbf{b}$, and the trained SFA and PCA instances. The output features $\mathbf{Y}$ are usually computed only during feature extraction (see Algorithm 2). Still, we keep them here to simplify the understanding of the signals involved.

5.3. Feature extraction

The algorithm for feature extraction is similar to the training algorithm, except that the parameters $\bar{\mathbf{x}}$, $\mathbf{W}_{\text{SFA}}$, $\bar{\mathbf{z}}$, $\mathbf{W}_{\text{PCA}}$, J, $\mathbf{Q}$, $\mathbf{R}$, $\mathbf{b}$ have already been

Algorithm 1 Training phase of iSFA

Require: $D > 0$: output dimension

procedure TRAIN($\mathbf{X}$) % $\mathbf{X} = (\mathbf{x_1}, \ldots, \mathbf{x_N})$: training samples
  $\forall n : \mathbf{x'_n} \leftarrow \mathbf{x_n} - \bar{\mathbf{x}}$ % $\bar{\mathbf{x}}$: average sample
  $\forall n : \mathbf{z_n} \leftarrow \mathbf{g}(\mathbf{x'_n})$
  $\mathbf{W}_{\text{SFA}}, \bar{\mathbf{z}} \leftarrow \text{SFA.TRAIN}(\mathbf{Z}, \text{OUTPUT_DIM} = J')$ % $\mathbf{Z} = (\mathbf{z_1}, \ldots, \mathbf{z_N})$: expanded samples
  $\forall n : \mathbf{s_n} \leftarrow \mathbf{W}_{\text{SFA}}(\mathbf{z_n} - \bar{\mathbf{z}})$
  $\forall n : \mathbf{s'_n} \leftarrow (\mathbf{s_{n1}}, \ldots, \mathbf{s_{nJ}})^T$ % Preserve the first J features with $\Delta < \Delta_T$
  $\forall n : \mathbf{a_n} \leftarrow \mathbf{M}\mathbf{s'_n} + \mathbf{b}$ % For $\mathbf{M}, \mathbf{b}$, such that $\mathbf{M}\mathbf{S'} + \mathbf{b} \approx \mathbf{X'}$
  $\forall n : \mathbf{y'_n} \leftarrow \mathbf{R}\mathbf{s'_n}$ % For $\mathbf{Q}\mathbf{R} = \mathbf{M}$, the QR decomposition of $\mathbf{M}$
  $\forall n : \mathbf{u_n} \leftarrow \mathbf{x'_n} - \mathbf{a_n}$
  $\mathbf{W}_{\text{PCA}} \leftarrow \text{PCA.TRAIN}(\mathbf{U}, \text{OUTPUT_DIM} = D - J)$
  $\forall n : \mathbf{h_n} \leftarrow \mathbf{W}_{\text{PCA}}\mathbf{u_n}$ % Only $D - J$ PCs are preserved
  $\forall n : \mathbf{y_n} = \mathbf{y'_n}|\mathbf{h_n}$ % '|' denotes vector concatenation
  return $\mathbf{Y} = (\mathbf{y_1}, \ldots, \mathbf{y_N})$, $\bar{\mathbf{x}}$, $\mathbf{W}_{\text{SFA}}$, $\bar{\mathbf{z}}$, $\mathbf{W}_{\text{PCA}}$, J, $\mathbf{Q}$, $\mathbf{R}$, $\mathbf{b}$
end procedure

learned from the training data. Algorithm 2 shows how a single input sample is processed, however, it can be easily and efficiently adapted to process multiple input samples by taking advantage of matrix operations.

One interesting property of iSFA is that the features are nonlinear w.r.t. the input data, both the slow and the reconstructive part. The slow part is nonlinear due to the expansion function. The residual data are nonlinear because it is computed using the (nonlinear) slow part and the centered data. The reconstructive part is thus only linear w.r.t the residual data but nonlinear w.r.t. the input data.

Algorithm 2 Feature extraction with iSFA

Require: D, $\bar{\mathbf{x}}$, $\mathbf{W}_{\text{SFA}}$, $\bar{\mathbf{z}}$, $\mathbf{W}_{\text{PCA}}$, J, $\mathbf{Q}$, $\mathbf{R}$, $\mathbf{b}$

procedure EXTRACT($\mathbf{x}$) % $\mathbf{x}$: a new sample
  $\mathbf{x'} \leftarrow \mathbf{x} - \bar{\mathbf{x}}$ % $\bar{\mathbf{x}}$: mean of the training data
  $\mathbf{z} \leftarrow \mathbf{g}(\mathbf{x'})$
  $\mathbf{s} \leftarrow \mathbf{W}_{\text{SFA}}(\mathbf{z_n} - \bar{\mathbf{z}})$ % Extract the first J slow features
  $\mathbf{s'} = (s_1, s_2, \ldots, s_J)$
  $\mathbf{y'} \leftarrow \mathbf{R}\mathbf{s'}$
  $\mathbf{a} \leftarrow \mathbf{Q}\mathbf{y'} + \mathbf{b}$
  $\mathbf{u} \leftarrow \mathbf{x'} - \mathbf{a}$
  $\mathbf{h} \leftarrow \mathbf{W}_{\text{PCA}}\mathbf{u}$
  $\mathbf{y} = \mathbf{y'}|\mathbf{h}$
  return $\mathbf{y}$
end procedure

5.4. Mixing and scaling of slow features

In the iSFA algorithm, the J-dimensional slow features $\mathbf{s}'_n$ are transformed into the scaled $\mathbf{y}'$ features. This transformation is necessary to make the amplitude of the slow features compatible with the amplitude of the PCA features, so that PCA processing of both sets of features together is possible in the next layers. The *QR scaling*, used by Algorithm 1, as well as a *sensitivity-based scaling*, are explained below.

Both methods ensure that the amplitude of the slow features is approximately equal to the reduction in the reconstruction error that they allow. In practice, a lower bound on the scales (not in the pseudo-code) ensures that all features have amplitudes > 0 even if they do not contribute to reconstruction.

A scaling method should ideally offer two key properties of PCA. (1) If one adds unit-variance noise to one of the principal components, the variance of the reconstruction error also increases by one unit. (2) If one adds independent noise to two or more principal components simultaneously, the variance of the reconstruction error increases additively.

Since $\mathbf{Q}$ is orthogonal, $\mathbf{y}'(t) \overset{(19,20,21)}{=} \mathbf{Q}^T(\mathbf{a}(t)-\mathbf{b})$ and $\mathbf{y}'$ is a rotation of the reconstructed data $\mathbf{a}$ (after centering). Thus, $\mathbf{y}'$ fulfills the two key properties of reconstruction of PCA w.r.t. $\mathbf{a}$.

One small drawback is that (21) mixes the slow features. Polynomial expansion functions combined with SFA are invariant to invertible linear transformations (e.g. SFA(QExp($\mathbf{Ux}$)) $\equiv$ SFA(QExp($\mathbf{x}$)), where $\mathbf{U}$ is any invertible matrix and QExp is the quadratic expansion). Thus, polynomial SFA can extract the same features from $\mathbf{s}'$ or $\mathbf{y}'$. However, other expansions do not have this property. One example of them is the function

$$0.8\text{Exp}(x_1, x_2, \ldots, x_I) \overset{\text{def}}{=} \\ (x_1, x_2, \ldots, x_I, |x_1|^{0.8}, |x_2|^{0.8}, \ldots, |x_I|^{0.8}), \tag{22}$$

which when combined with SFA is invariant to scalings of the input data but not to their mixing, i.e., SFA(0.8Exp($\mathbf{\Lambda x}$)) $\equiv$ SFA(0.8Exp($\mathbf{x}$)), where $\mathbf{\Lambda}$ is a diagonal matrix with diagonal elements $\lambda_i \neq 0$, but SFA(0.8Exp($\mathbf{Ux}$)) $\not\equiv$ SFA(0.8Exp($\mathbf{x}$)) in general. The 0.8Exp expansion has been motivated by a model where the slow features are noisy harmonics of increasing frequency of a hidden parameter and it should be applied to the slow features directly. Thus, the mixture of the slow features would break the assumed model, which might compromise slowness extraction in the next layers.

Technically, such mixing could be reverted in the next layer (e.g., by an additional application of linear SFA before the expansion), but this would add unnecessary complexity. For this reason, besides the QR scaling, we propose a second scaling method. The sensitivity based scaling scales the slow features without mixing them, as follows.

$$\mathbf{y}' = \mathbf{\Lambda s}', \tag{23}$$

where $\mathbf{\Lambda}$ is a diagonal matrix with diagonal elements $\lambda_j \stackrel{\text{def}}{=} ||\mathbf{M}_j||_2$ (the L_2-norm of the j-th column vector of $\mathbf{M}$). Therefore,

$$\mathbf{a}(t) \stackrel{(23)}{=} \mathbf{M}\mathbf{\Lambda}^{-1}\mathbf{y}'(t) + \mathbf{b}\,. \tag{24}$$

Clearly, the transformation (23) does not mix the slow features, it only scales them. From the two key reconstruction properties of PCA mentioned above (adding noise of certain variance to either one or many features increases the variance of the reconstruction error by the same amount), the first one (noise on a single feature) is fulfilled, because the columns of $\mathbf{M}\mathbf{\Lambda}^{-1}$ have unit norm. The second one is not fulfilled, because $\mathbf{M}\mathbf{\Lambda}^{-1}$ is in general not orthogonal.

On a first sight, it seems like multiple-view learning might be an alternative for the scaling methods used here. However, it is actually solves a different problem. Moreover, it is not suitable for this task because it would mix the slow and reconstructive parts, and might be expensive computationally (e.g., [32]) with cubic complexity in N.

5.5. Input reconstruction for iSFA

iSFA allows a linear approximation of the input (linear input reconstruction) even though all the features computed are nonlinear. In contrast, standard SFA does not have a standard input reconstruction method, although various gradient-descent and vector-quantization methods have been tried (e.g., [30]) with limited success.

In the reconstruction algorithm, $\mathbf{a}$ (the contribution of the slow part to the centered data) is approximated as $\tilde{\mathbf{a}} = \mathbf{Q}\mathbf{y}' + \mathbf{b}$, where tilde denotes an approximation. Then, $\mathbf{u}$ (the residual vector) is approximated as $\tilde{\mathbf{u}} = \mathbf{W}_{\text{PCA}}^T\mathbf{h}$. The reconstructed sample is then $\tilde{\mathbf{x}} = \tilde{\mathbf{a}} + \tilde{\mathbf{u}} + \bar{\mathbf{x}}$. See Algorithm 3 for details.

Algorithm 3 Linear input reconstruction for iSFA

Require: D, $\bar{\mathbf{x}}$, $\mathbf{W}_{\text{PCA}}$, J, $\mathbf{Q}$, $\mathbf{R}$, $\mathbf{b}$

1: **procedure** LINEAR-RECONSTRUCTION($\mathbf{y}$)
2: $\quad\mathbf{y}' \leftarrow (y_1, \ldots, y_J)$ % Slow part
3: $\quad\mathbf{h} \leftarrow (y_{J+1}, \ldots, y_D)$ % Reconstructive part
4: $\quad\tilde{\mathbf{x}} \leftarrow (\mathbf{Q}\mathbf{y}' + \mathbf{b}) + \mathbf{W}_{\text{PCA}}^T\mathbf{h} + \bar{\mathbf{x}}$ % $\tilde{\mathbf{a}} + \tilde{\mathbf{u}} + \bar{\mathbf{x}}$
5: $\quad$**return** $\tilde{\mathbf{x}}$
6: **end procedure**

The linear reconstruction algorithm has interesting properties. Firstly, it is simpler than the feature extraction algorithm, because the nonlinear expansion and $\mathbf{W}_{\text{SFA}}$ are not needed. Secondly, it has a low computational complexity, because it consists of only two matrix-vector multiplications and a couple of vector additions, all of them I, J, $D - J$ or D-dimensional.

Linear reconstruction is simple and effective for iSFA. We now describe a nonlinear reconstruction algorithm. Assume $\mathbf{y}$ is the iSFA feature representation of a sample $\mathbf{x}$, which we denote as $\mathbf{y} = \text{iSFA}(\mathbf{x})$. Since $\mathbf{x}$ is unknown, the

reconstruction error cannot be computed directly. However, one can indirectly measure the accuracy of a particular reconstruction $\tilde{\mathbf{x}}$ by means of the feature error, which is defined here as $e_{\text{feat}} \stackrel{\text{def}}{=} ||\mathbf{y} - \text{iSFA}(\tilde{\mathbf{x}})||$. This feature error can be minimized for $\tilde{\mathbf{x}}$ using the function iSFA$(\cdot)$ as a black box and gradient descent or other generic nonlinear minimization algorithms. Frequently, such algorithms require a first approximation, which can be very conveniently provided by the linear reconstruction algorithm.

Although nonlinear reconstruction methods might result in higher reconstruction accuracy, they are typically more expensive computationally. Moreover, in the discussion we explain why minimizing e_{feat} does not necessarily improve the reconstruction error unless other aspects are considered.

5.6. Some remarks on iSFA

Clearly, the computational complexity of iSFA is at least that of SFA, because iSFA consists of SFA and a few additional computations. However, none of these additional computations is done on the expanded I'-dimensional data but at most on I or D-dimensional data (e.g., PCA is applied to I-dimensional data, and the QR decomposition is applied to a $I \times I$-matrix resulting in $\mathcal{O}(NI^2 + I^3)$ and $\mathcal{O}(I^3)$ complexity, respectively). Therefore, iSFA is slightly slower than SFA but it has the same complexity order. Practical experiments (Section 6) confirm this.

The presentation above focuses on iSFA. To obtain information-preserving GSFA (iGSFA) one only needs to substitute SFA by GSFA inside the iSFA algorithm and provide GSFA with the corresponding training graph during the training phase. Notice that the GSFA features have weighted zero-mean instead of the simple (unweighted) zero-mean enforced by SFA. This small difference has already been compensated by the vector $\mathbf{b}$.

We have implemented the iGSFA node (including iSFA) in python. In the next few months, we plan to make the node public by integrating it to the MDP toolkit [33].

6. Experimental evaluation

In this section, we evaluate HiGSFA using the problem of age estimation from human face photographs. HiGSFA is employed instead of HiSFA, because we use the labels to boost estimation accuracy. However, due to their close connection, many aspects of the evaluation also extend to HiSFA.

First, age estimation and previous work are introduced. Then, the preprocessing of the images is described. Afterwards, a new training graph for learning age, race and gender is proposed. Then, an HiGSFA network is described and evaluated according to three criteria: feature slowness (compared with HGSFA), age estimation error (compared with state-of-the-art algorithms), and linear reconstruction error (compared with PCA).

6.1. Age estimation and previous work on this problem

The study of age estimation from photographs is relatively recent and has useful applications for human-computer interaction, group-targeted advertisement, demographics, face recognition, control of age-related policies, and security. However, age estimation is a challenging task probably because different persons experience facial aging differently depending on factors such as their gender, race, life style, nutrition, health, exposure to the weather, use of creams or cosmetics, operations, accidents, and even psychological traits.

Two types of age have been distinguished. The real (ground-truth) age, which is the chronological age of the person, and the apparent age, which is the age conveyed solely by the information present in the image. Clearly, an algorithm cannot determine the real age exactly, at most it might determine the apparent age.

Different methods for age estimation have been proposed. For a more comprehensive literature review we refer to [34] and [35]. In [36], aging pattern subspace (AGES) has been proposed, which is based on temporal sequences of images of individual persons, the so-called aging patterns. The images are represented through an appearance model that combines geometric and texture information. In their system, a subspace is constructed for each aging pattern. Given a new image, the subspace providing the best possible reconstruction is found. Then, the position of the image within the aging pattern is determined.

Guo et al. [37] have proposed the use of bio-inspired features (BIF). Their architecture consists of two-layers, in which the units of the first layer compute Gabor functions inspired by simple cells, whereas the units of the second layer compute a standard-deviation operation inspired by complex cells. Then, PCA is applied to reduce the dimensionality of the data to fewer than 1,000 features. Finally, an SVM or SVR provides the final age estimate.

The first SFA architecture for age estimation was a four-layer HSFA network that was applied to the images directly without prior feature extraction [11]. The input images were synthetic and created using special software for 3D-face modeling. However, the complexity of the face model was probably too simple, which allowed linear SFA to achieve good performance, and left open the question of whether SFA could also be successful on real photographs.

Race and gender are two factors that influence the accuracy of age estimation (e.g., [38, 39]). This idea is exploited by the system in [40], where the faces are first classified according to race and gender, and age is then estimated in the particular race/gender group. Other algorithms allow the estimation of age, race and gender simultaneously. Guo et al. [41] proposed the use of kernel partial least squares regression (KPLS) on top of the BIF features.

More recently, various methods based on canonical correlation analysis (CCA) have been proposed by Guo et al. [42], particularly regularized kernel CCA (rKCCA), on top of BIF features in a framework for the joint estimation of age, race and gender.

The use of a multi-scale convolutional neural network (MCNN) trained on 23 48×48 image patches has been proposed in [43]. Each patch has one out of four different scales and is centered on a particular facial landmark.

6.2. Image database and image pre-processing

The MORPH-II database (i.e. MORPH, Album 2) [44] is a large database available for a symbolic fee and suitable for age estimation. It contains 55 134 images of about 13 000 different persons with ages ranging from 16 to 77 years. The images were taken under partially controlled conditions (e.g. frontal pose, absence of glasses, good image quality, absence of strong shadows), and include variations in head pose (e.g. tilt angle) and expression. The database includes annotations stating the age of the persons, their gender (M or F), "race": "black" (B), "white" (W), "asian" (A), "hispanic" (H), and "other" (O), and the coordinates of the eyes. The procedure used to assign the race label is unknown to us. Most of the images are of black (77%) or white races (19%), probably making it more difficult to generalize to other races, such as asian. We chose this database because of its large number of images.

We follow the evaluation method used in [42] and many other works. The input images are partitioned in 3 disjoint sets S_1 and S_2 of 10 530 images, and S_3 of 34 074 images. The racial and gender composition of S_1 and S_2 is the same: they have about 3 times more images of males than females and the same number of white and black people. Other races are omitted. More exactly, $|MB| = |MW| = 3980$, $|FB| = |FW| = 1285$. The remaining images constitute the set S_3, which is composed as follows: $|MB| = 28\,872$, $|FB| = 3187$, $|MW| = 1$, $|FW| = 28$, $|MA| = 141$, $|MH| = 1667$, $|MO| = 44$, $|FA| = 13$, $|FH| = 102$ and $|FO| = 19$. Training and testing are done twice, using either S_1 and S_1-test $\stackrel{\text{def}}{=} S_2 + S_3$ or S_2 and S_2-test $\stackrel{\text{def}}{=} S_1 + S_3$.

We pre-process the input images in two steps: pose normalization and face sampling (Figure 2). The pose-normalization step fixes the position of the eyes ensuring that: a) the eye line is horizontal, b) the inter-eye distance is constant, and c) the output resolution is 256×260 pixels. After pose normalization, face sampling keeps the head area only, enhances the contrast, and scales down the image to 96×96 pixels. Typically the chin, forehead, and some hair are visible in the resulting images.

We define a DR-dataset to train HiGSFA (or the DR algorithm), an S-dataset to train the supervised step on top of HiGSFA (a Gaussian classifier), and a T-dataset for testing. The DR and S-datasets are created with the training images (either S_1 or S_2), and the T-dataset with the corresponding test images, either S_1-test or S_2-test.

The images of the DR and S-datasets go through a distortion step during face sampling, which include a small random translation of max ±1.4 pixels, a rotation of max ±2 degrees, and a rescaling of ±4%, as well as small fluctuations in the average color and contrast. The exact transformations are distributed uniformly in their corresponding intervals. Although the distortions are frequently imperceptible, they teach HiGSFA to become invariant to small errors during image normalization and are necessary due to its feature specificity to improve generalization to test data. Other algorithms that use particular structures (e.g., convolutional layers, max pooling) or BIF features are mostly invariant to such small transformations by construction (e.g., [42]).

Distortions allow us to increase the number of images used for training. The images of the DR-dataset are distorted 22 times with different random distortions and those of the S-dataset 3 times, resulting in 231 660 and 31 590 images, respectively. The images of the T-dataset are not distorted and used only once.

6.3. Efficient training graphs for learning multiple-labels

GSFA would be less attractive without its efficient training graphs, which include a clustered graph (classification) and a serial graph (regression) with a training complexity of $\mathcal{O}(NI^2 + I^3)$. However, up to now efficient graphs have only been defined for a single (categorical or numerical) label. We briefly recall the clustered and serial graphs [10], and then propose an efficient graph for 3 labels.

Clustered training graph. This graph generates features useful for classification that are equivalent of those of FDA, see [45] (also compare [46] and [22]), and is illustrated in Figure 7. The optimization problem associated with this graph explicitly demands that samples from the same class should be typically mapped to similar outputs. If C is the number of classes, $C-1$ output features can be extracted and given to a standard classifier to compute the final class estimate.

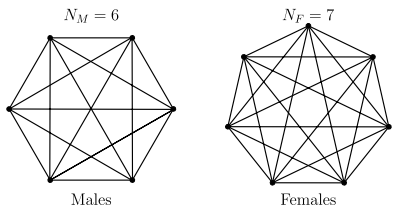

Figure 7: Illustration of a *clustered* training graph for gender classification with 6 images of males and 7 of females. Each vertex represents an image and edges represent transitions. Pairs of images $\mathbf{x}(n)$ and $\mathbf{x}(n')$, with $n \neq n'$, that belong to the same class s (either F or M) are connected with an edge weight $\gamma_{n,n'} = 1/(N_s - 1)$, where N_s is the number of images in class s. This results in 2 fully connected subgraphs. Images of different genders are not connected. The weight of all vertices is equal to one. For the actual experiments, we use $N_F = 2570r$ and $N_M = 7960r$, for $r \stackrel{\text{def}}{=} 22$. The graph for race (B and W) is similar with $N_B = N_W = 5265r$.

Serial training graph. This graph has consistently given good results for different regression problems. Other efficient training graphs for regression are the sliding window and mixed graphs [10]. After GSFA, a complementary explicit regression step on a few features solves the original regression problem.

The serial graph is constructed by ordering the samples by increasing label. Then, the samples are partitioned into L groups of size $N_g = N/L$. Each group has a representative label $\in \{\ell_1, \ldots, \ell_L\}$, where $\ell_1 < \ell_2 < \cdots < \ell_L$, see Figure 8. Edges connect all pairs of samples from two consecutive groups with representative labels (ℓ_l and ℓ_{l+1}). Thus, all connections are inter-group, no

intra-group connections are present. Notice that since any two vertices of the same group are adjacent to exactly the same neighbors, they are likely to be mapped to similar outputs by GSFA.

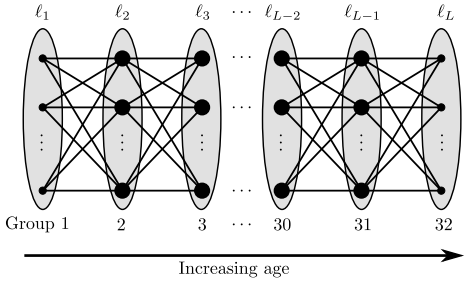

Figure 8: Illustration of a serial training graph for age estimation. The training images are first ordered by increasing age and then grouped into $L = 32$ groups of $N_g = 7238$ samples each. Each dot represents an image, edges represent connections, and ovals represent the groups. The images of the first and last group have weight 1 and the remaining images have weight 2 (image weights represented by smaller/bigger dots). The weight of all edges is 1.

The theory of unrestricted GSFA predicts that mapping the slowest feature may suffice to solve the regression problem accurately (except for a discretization error). This requires that the slowest hidden parameter of the data is a strictly monotonic function of the label. In practice, however, mapping a few slow features frequently gives better results. Different approaches for implementing this mapping have been proposed [10].

Efficient graph for age, race and gender estimation. Based on a recent analysis of the optimal free responses of GSFA [12] (i.e., the slowest features possible when the feature space is unlimited), we propose here the combination of K efficient graphs that learn one label each into a single graph that learns K labels. To compute the combined graph, one must only add the vertex and edge weights. For the approach to be mathematically sound, we require 3 conditions: (1) all graphs have the same samples and are consistent, (2) all graphs have the same (or proportional) node weights, and (3) optimal free responses that are slow in one graph ($\Delta < 2.0$) should not be fast ($\Delta > 2.0$) in any other graph. These requirements guarantee that the slowest optimal free responses of the combined graph span the slowest optimal free responses of the original graphs.

For example, one can combine a clustered graph for gender (M or F) estimation and another for race (B or W). The first two features of the resulting graph are then enough for gender *and* race classification. Alternatively, one could create a clustered graph with four classes (MB, MW, FB, FW), so that 3 features would be needed for classification instead of 2. However, this would be impractical for larger numbers of classes. For example, if the original numbers of classes were $C_1 = 10$ and $C_2 = 12$, one would need to extract

$C_1C_2 - 1 = 119$ features, whereas in the proposed graph, one would only need to preserve $(C_1 - 1) + (C_2 - 1) = 20$ features.

To learn the age, race and gender labels, we propose a graph G_3 that combines a serial graph for age estimation, a clustered graph for gender, and a clustered graph for race. The serial graph has node weights not quite the same as in the clustered graphs, but this might not affect the accuracy of the combined graph significantly. For comparison, we also use a serial graph G_1 that only learns age.

We use the first 4 to 7 features extracted from the S-dataset to train 3 separate Gaussian classifiers (GC). For race only 2 classes are considered (B and W), and for gender only M and F. For age, the images are partitioned in 39 classes of increasing ages. The classes have average ages of $\{16.6, 17.6, 18.4, \ldots, 52.8, 57.8\}$ years. To compute these average ages, as well as to order the samples by age in the serial graph, we use the age of the persons with a day resolution (e.g., 25.216 years). However, for the evaluation, we use integer ground-truth labels and integer estimates.

After the GC has been trained, the final age estimation (on the T-dataset) is computed using class membership probabilities. Let $\mathbf{P}(C_{\tilde{\ell}_l}|\mathbf{y})$ be the estimated class probability that the input sample $\mathbf{x}$ with feature representation $\mathbf{y} = \mathbf{g}(\mathbf{x})$ belongs to the group with average label $\tilde{\ell}_l$. Then, the estimated age is

$$\ell \stackrel{\text{def}}{=} \sum_{l=1}^{39} \tilde{\ell}_l \cdot \mathbf{P}(C_{\tilde{\ell}_l}|\mathbf{y})\,, \tag{25}$$

where a final step keeps the integer part of ℓ. Equation (25) is particularly suited to minimize the root mean squared error (RMSE). Although it incurs in an error due to the discretization of the labels, the soft nature of the estimation has provided good accuracy and robustness.

6.4. Evaluated algorithms

Besides evaluating HiGSFA, we resort to HGSFA, PCA, and state-of-the-art age-estimation algorithms for comparison purposes. iGSFA and GSFA are not used directly, but HiGSFA and HGSFA are used to take advantage of the benefits of hierarchical processing.

The structure of the HiGSFA and HGSFA networks is described in Table 2. In both networks, the nodes are composed only of iGSFA or GSFA, with various types of nonlinear expansion functions. Only in the first layer, PCA is applied to the pixel data prior to iGSFA/GSFA, preserving 20 out of 36 principal components. To scale the slow features we use the sensitivity method of Section 5.4. The hyper-parameters have been hand-tuned to achieve best accuracy on age estimation using educated guesses, random sets S_1, S_2 and S_3 different to those used for the evaluation, and fewer image multiplicities to speed up the process.

The HGSFA/HiGSFA networks differ in several aspects from SFA networks used in the literature (e.g., [27]). For example, to improve feature specificity at the lowest layers, no weight sharing was used. Moreover, the input to the nodes

(fan-in) originates mostly from the output of 3 nodes in the preceding layer (3×1 or 1×3). Such a small fan-in reduces the computational cost because the input dimensionality is minimized. It also results in networks with 10 layers, potentiating the accumulation of nonlinearity across the network.

The expansion functions are a mixture of different nonlinear functions on subsets of the input vectors, including: (1) The identity function $\mathrm{I}(\mathbf{x}) = \mathbf{x}$. (2) The quadratic terms $\mathrm{QT}(\mathbf{x}) \stackrel{\text{def}}{=} \{x_i x_j\}_{i,j=1}^{N}$. (3) A normalized version of QT: $\mathrm{QN}(\mathbf{x}) \stackrel{\text{def}}{=} \{\frac{1}{1+||\mathbf{x}||^2} x_i x_j\}_{i,j=1}^{N}$. (4) The terms $0.8\mathrm{ET}(\mathbf{x}) \stackrel{\text{def}}{=} \{|x_i|^{0.8}\}_{i=1}^{N}$, which are useful to improve generalization and resistance against outliers [29]. (5) The function $\textsc{max}2(\mathbf{x}) \stackrel{\text{def}}{=} \{\max(x_i, x_{i+1})\}_{i=1}^{N-1}$. We propose this function inspired by the state-of-the-art algorithms on age estimation, which include max pooling or some variation of it. For example, the nonlinear expansion of the first layer of the HiGSFA network is $\mathrm{I}(x_1, \ldots, x_{18}) \,|\, 0.8\mathrm{ET}(x_1, \ldots, x_{15}) \,|\, \textsc{max}2(x_1, \ldots, x_{17}) \,|\, \mathrm{QT}(x_1, \ldots, x_{10})$, where $|$ indicates vector concatenation. The details of the expansions used in the remaining layers are available upon request.

Layer	size	node fan-in	stride	output dim. per node	
				HGSFA	HiGSFA
0	96×96 pixels	—	—	—	—
1	31×31 nodes	6×6	3×3	14	18
2	15×31 nodes	3×1	2×1	20	27
3	15×15 nodes	1×3	1×2	27	37
4	7×15 nodes	3×1	2×1	49	66
5	7×7 nodes	1×3	1×2	60	79
6	3×7 nodes	3×1	2×1	61	88
7	3×3 nodes	1×3	1×2	65	88
8	1×3 nodes	3×1	1×1	65	93
9	1×1 nodes	1×3	—	66	95
10	1×1 nodes	1×1	—	75	75

Table 2: Description of the HiGSFA and HGSFA networks. The two networks have the same number of nodes and general structure, but they differ in the type of nodes and in the number of features these preserve. Layer 0 denotes the input image, whereas layer 10 is the top node. The size of the slow part in layers 1 and 2 does not depend on Δ_T but is fixed to 3 or 4 features, resp.

The parameter Δ_T of layers 3 to 10 was set to 1.96. Δ_T was not used in layers 1 and 2, where the number of slow features was fixed to 3 and 4, resp. The number of features given to the supervised algorithm, shown in Table 3, was tuned for each DR algorithm and supervised problem.

Since the data dimensionality allows it, we used PCA directly (contrary to hierarchical PCA) to provide more accurate principal components and smaller reconstruction errors.

6.5. Experimental results

We show now the results of HiGSFA, HGSFA and PCA on the three evaluation criteria. Results are reported as $a \pm b$, where a is the average over the

Algorithm	Age	Race	Gender
HiGSFA (G_3)	5	6	4
HGSFA (G_3)	5	5	7
PCA	54	54	60

Table 3: Number of output features given to the supervised step (a Gaussian classifier).

test images (S_1-test and S_2-test), and b is the standard error of the mean (i.e., half the absolute difference).

6.5.1. Feature slowness

The weighted Δ values of GSFA (Equation 1), which we denote here as $\Delta_j^{\mathrm{DR},G_3}$, depend on the graph G_3, which in turn depends on the training data and their labels. To measure slowness (or rather fastness) of test data, we compute standard Δ values on the images ordered by increasing age label, $\Delta_j^{\mathrm{T,lin}} \stackrel{\mathrm{def}}{=} \frac{1}{N-1}\sum_n (y_j(n+1) - y_j(n))^2$, which is equivalent to using a linear graph. In all cases, the scale of the features is normalized to variance 1 before computing the Δ values to prevent alterations due to the feature scaling method.

Table 4 shows $\Delta_{1,2,3}^{\mathrm{DR},G_3}$ (resp. $\Delta_{1,2,3}^{\mathrm{T,lin}}$), that is, the Δ values of the three slowest features extracted from the DR-dataset (resp. T-dataset) using the graph G_3 (resp. a linear graph).

	PCA	HGSFA (G_3)	HiGSFA (G_3)
$\Delta_1^{\mathrm{DR},G_3}$	–	1.23	1.17
$\Delta_2^{\mathrm{DR},G_3}$	–	1.46	1.38
$\Delta_3^{\mathrm{DR},G_3}$	–	1.56	1.53
$\Delta_1^{\mathrm{T,lin}}$	1.99	0.45	0.38
$\Delta_2^{\mathrm{T,lin}}$	1.93	1.12	0.99
$\Delta_3^{\mathrm{T,lin}}$	1.99	1.90	1.90

Table 4: Average delta values of the first three features extracted by PCA, HGSFA and HiGSFA on training and test data (the smaller the better). The first feature extracted is the most stable according to the age-ordered linear graph, which indicates that this is the main feature coding age. For comparison, the Δ value of unit-variance i.i.d. noise is 2.0.

Table 4 shows that HiGSFA outperforms HGSFA in slowness maximization. The $\Delta^{\mathrm{T,lin}}$ values of the PCA features are larger, which is not surprising, because PCA does not optimize for slowness. Since Δ^{DR,G_3} and $\Delta^{\mathrm{T,lin}}$ are computed from different graphs, they should not be compared with each other. $\Delta^{\mathrm{T,lin}}$ considers transitions between images with the same or very similar ages but arbitrary race and gender. Δ^{DR,G_3} only considers transitions between images having at least one of a) the same gender, b) the same race, or c) different but consecutive age groups.

6.5.2. Age estimation error

Some real-life applications only need a coarse categorization of age in broad age groups. However, other applications benefit from a more precise estimation, making it convenient to treat age estimation as a regression problem requiring a concrete numerical estimation, usually expressed as an integer number of years.

We use three metrics to measure age estimation accuracy: (1) The mean absolute error (MAE), see [36], which is the most frequent metric for age estimation, (2) the root mean squared error (RMSE), which is a common loss function for regression in Machine Learning, and (3) cumulative scores (CSs), see [36], which indicate the fraction of the images that have an estimation error below a given threshold. For instance, $CS(5)$ is the fraction of estimates (e.g., expressed as a percentage) having an error of at most 5 years w.r.t. the real age. We include CSs at various thresholds, facilitating future comparisons with other methods. Although the RMSE is sensitive to outliers and has almost not been used in the literature on age estimation, some applications might benefit from its stronger penalization of larger estimation errors. The accuracies are summarized in Table 5.

Algorithm	Age (MAE)	Age (RMSE)	Race (CR)	Gender (CR)
BIF+3Step [40]	4.45 ± 0.01	—	$98.80\% \pm 0.04$	$97.84\% \pm 0.16$
BIF+KPLS [41]	4.18 ± 0.03	—	$98.85\% \pm 0.05$	$98.20\% \pm 0.00$
BIF+rKCCA [42]	3.98 ± 0.03	—	$99.00\% \pm 0.00$	$\mathbf{98.45\% \pm 0.05}$
BIF+rKCCA+SVM [42]	3.92 ± 0.02	—	—	—
baseline CNN [43]	4.60 ± 0.05	—	—	—
MCNN* no align [43]	3.79 ± 0.09	—	—	—
MCNN* only age [43]	3.63 ± 0.00	—	—	—
MCNN* [43]	3.63 ± 0.09	—	$98.6\% \pm 0.05$	$97.9\% \pm 0.1$
PCA (control)	6.804 ± 0.007	8.888 ± 0.000	$96.75\% \pm 0.06$	$91.54\% \pm 0.24$
HGSFA (G_3) (control)	3.921 ± 0.018	5.148 ± 0.049	$98.60\% \pm 0.08$	$96.40\% \pm 0.12$
HiGSFA (G_1) (control)	3.605 ± 0.001	4.690 ± 0.000	—	—
HiGSFA (G_3) (proposed)	$\mathbf{3.497 \pm 0.008}$	$\mathbf{4.583 \pm 0.000}$	$\mathbf{99.15\% \pm 0.01}$	$97.70\% \pm 0.01$
Chance level	9.33	10.95	87.58%	86.73%

Table 5: Accuracy in years of state-of-the-art algorithms for age estimation on the MORPH-II database (test data). Classification rates (CR) for race and gender estimation are also provided. The chance level is the best possible performance when the estimation is constant. *A mistake in the evaluation protocol of MCNN made their training and test data not disjoint, thus the actual accuracy might differ, see `http://www.cbsr.ia.ac.cn/users/dyi/agr.html`.

The MAE of HGSFA was 3.921 years, which is better than that of BIF+3Step, BIF+KPLS and BIF+rKCCA, similar to BIF+rKCCA+SVM, and worse than the MCNNs (except the baseline). However, the MAE of HiGSFA is only 3.497 years, which is better than all previous algorithms we have found published. In contrast, with an MAE of 6.804 years PCA has the largest MAE. Detailed cumulative scores for HiGSFA and HGSFA are provided in Table 6.

The RMSE of HGSFA on test data is 5.148 years, HiGSFA yielded an RMSE of 4.583 years, and PCA an RMSE of 8.888 years. The RMSE of other ap-

proaches does not seem to be available.

The poor accuracy of PCA for age estimation is not surprising, because principal components might lose wrinkles, skin imperfections, and other information that could reveal age in adult persons, and it might also be explained because principal components are too unstructured to be properly untangled by the Soft GC method, in contrast to slow features, which have a very specific and simple structure.

Algorithm	cs(0)	cs(1)	cs(2)	cs(3)	cs(4)	cs(5)	cs(6)	cs(7)	cs(8)	cs(9)	cs(10)
HGSFA	8.80	26.42	42.41	55.80	66.38	74.86	81.31	86.37	90.12	92.93	94.96
HiGSFA	9.87	29.16	46.23	60.14	71.04	79.56	85.70	90.04	93.12	95.45	96.92

	cs(11)	cs(12)	cs(14)	cs(16)	cs(18)	cs(20)	cs(22)	cs(24)	cs(26)	cs(28)	cs(30)
HGSFA	96.43	97.42	98.67	99.34	99.69	99.85	99.92	99.97	99.97	99.98	99.99
HiGSFA	97.91	98.56	99.34	99.70	99.84	99.91	99.94	99.96	99.97	99.98	99.99

Table 6: Percentile cumulative scores (the larger the better) for various maximum allowed errors ranging from 0 to 30 years.

The behavior of the estimation errors as a function of the real age is plotted in Figure 9. On average, older persons are estimated much younger than they really are. This is in part due to the small number of older persons in the database, and because the oldest class used in the supervised step has an average of about 58 years, making this the largest age that can be estimated. The MAE is surprisingly low for persons below 45 years. For 19-year-old persons, the MAE is only 2.253 years.

6.5.3. Reconstruction error

A reconstruction error provides an indication of how much information the output features retain from the original input. In order to compute it, we assume a linear global model for input reconstruction.

Let $\mathbf{X}$ be the input data and $\mathbf{Y}$ the corresponding set of extracted features. A matrix $\mathbf{D}$ and a vector $\mathbf{c}$ are learned from the DR-dataset using linear regression such that $\hat{\mathbf{X}} \stackrel{\text{def}}{=} \mathbf{DY} + \mathbf{c1}^T$ approximates $\mathbf{X}$ as closely as possible, where $\mathbf{1}$ is a vector N of ones. $\hat{\mathbf{X}}$ is a matrix containing the reconstructed samples (i.e. $\hat{\mathbf{x}}_\mathbf{n} \stackrel{\text{def}}{=} \mathbf{D}\mathbf{y}_\mathbf{n} + \mathbf{c}$ is the reconstruction of the input $\mathbf{x}_\mathbf{n}$ given its feature representation $\mathbf{y}_\mathbf{n}$). Figure 2 provides examples of face reconstructions from different features.

Since the model is linear and global the output features are mapped to the input linearly. For PCA this gives the same result as the usual multiplication with the transposed projection matrix plus image average. An alternative (local) approach for HiGSFA would be to use the pseudo-inversion algorithm to perform reconstruction from the top of the network to the bottom, one node at a time.

The normalized reconstruction error, computed from the T-dataset, is then

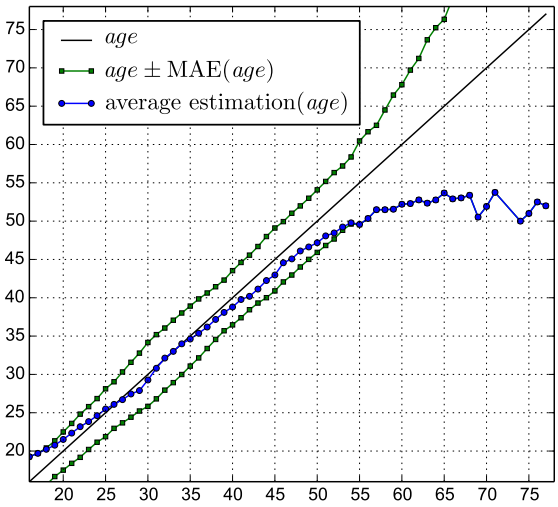

Figure 9: The average age estimate is plotted as a function of the real age. The MAE is also computed as a function of the real age and plotted as $age \pm \mathrm{MAE}(age)$.

defined as

$$e_{\mathrm{rec}} \stackrel{\mathrm{def}}{=} \frac{\sum_{n=1}^{N} ||(\mathbf{x_n} - \hat{\mathbf{x}}_\mathbf{n})||^2}{\sum_{n=1}^{N} ||(\mathbf{x_n} - \bar{\mathbf{x}})||^2}\,, \tag{26}$$

which is the ratio between the energy of the reconstruction error and the variance of the test data except for a factor $N/N-1$.

	Chance level	HGSFA	HiGSFA	PCA
e_{rec}	1.0	0.818	0.338	**0.201**

Table 7: Reconstruction errors on test data using 75 features and various algorithms.

The reconstruction errors of HGSFA, HiGSFA and PCA using 75 features are given in Table 7. The largest reconstruction error results from the constant reconstruction $\bar{\mathbf{x}}$ (chance level). As expected, HGSFA does slightly better than chance level, but worse than HiGSFA, which is closer to PCA. PCA yields the best possible features for the given linear global reconstruction method, and is better than HiGSFA by 0.127. For HiGSFA, from the 75 output features, 8 of them are slow features (slow part), and the remaining 67 are reconstructive. Using 67 features, PCA yields a reconstruction error of 0.211.

6.5.4. HiGSFA network with HGSFA hyper-parameters

In the experiments above, the hyper-parameters of the HiGSFA and HGSFA networks (e.g., nonlinear expansion functions, output dimensionalities) were tuned separately. To verify that the performance of HiGSFA is better than that

of HGSFA not simply due to different hyper-parameters, we evaluated the performance of an HiGSFA network using the hyper-parameters of the HGSFA network (the only difference is the use of iGSFA nodes instead of GSFA nodes). The hyper-parameter Δ_T, not present in HGSFA, was set as in the tuned HiGSFA network. As expected, the performance of HiGSFA was affected: The MAE increased to 3.72 years, and the RMSE increased to 4.80 years. The reconstruction error improved slightly to 0.322. Although the suboptimal hyper-parameters affected HiGSFA, it was still clearly superior to HGSFA.

6.5.5. Sensitivity to the delta threshold Δ_T

We have evaluated the influence of Δ_T on estimation accuracy and numerical stability, by testing different values for Δ_T. For simplicity, we used the same Δ_T from layers 3 to 10 in this experiment (Δ_T is not used in layers 1 and 2, where the number of features in the slow part is constant, 3 and 4 features, respectively). The performance of the algorithm as a function of Δ_T is shown in Table 8. The Δ_T yielding minimum MAE and used in the optimized architecture is 1.96.

Δ_T	Age MAE	Age RMSE	e_{rec}	Race CR	Gender CR	#Features L3
1.92	3.506	4.583	0.330	99.15	97.58	2.87
1.94	3.499	4.583	0.333	99.15	97.64	3.22
1.96	3.497	4.583	0.338	99.15	97.00	3.66
1.98	3.530	4.637	0.359	99.09	94.72	4.14

Table 8: Performance of HiGSFA on the MORPH-II database using different Δ_T (default value is $\Delta_T = 1.96$). The results reported are the age estimation errors (MAE and RMSE), the reconstruction error e_{rec}, the percentile classification rate for race and gender, and the average number of features with $\Delta < \Delta_T$ in the nodes of the third HiGSFA layer. All error measures were computed on test data.

The average number of slow features in the third layer changed moderately depending on Δ_T, ranging from 2.87 to 4.14 features, and the differences in the final error measures were small. This shows that the parameter Δ_T is not critical and can be easily tuned.

6.5.6. Evaluation on the FG-NET database

We investigated the capability of HiGSFA to generalize to a different database. We used the HiGSFA (G_3) network trained with images of the MORPH-II database (either with the set S_1 or S_2), and tested it using images of the FG-NET database [47]. The test images were taken under uncontrolled conditions (e.g., many are not frontal) and we excluded those out of the original age range of 16 to 77 years.

For age estimation, the MAE was 7.32 $\pm$ 0.08 years and the RMSE 9.51 $\pm$ 0.13 years (using 4 features for the supervised step). For gender and race estimation, the classification rates (5 features) are 80.85% $\pm$ 0.95% and 89.24% $\pm$ 1.06%, resp. We assumed white race for all the test persons.

The most comparable cross-database experiment known to us is a system trained on a large database of images from the internet and tested on FG-Net [48]. If the same age range as above is used, the MAE is approximately 8.29 years.

6.5.7. Alternative evaluation protocol

The protocol used to create the training and test data (based on the sets S_1, S_2 and S_3) [42] has been frequently used for age estimation. However, it has a few disadvantages. Thus, we also adopt the "leave one person out" (LOPO) [49, 50] protocol. For efficiency reasons, we leave 2000 persons out of the training set instead of one. That is, the test data are created with the images of 2000 persons chosen at random (about 8000 images), whereas the training data is simply the remaining images (about 45 000 images).

This protocol has the following properties in contrast to the one used in [42]: (a) The distribution of the training and test images is the same on average: $|M| = 85\%$, $|F| = 15\%$, $|B| = 77\%$, $|W| = 19\%$, $|H| = 3\%$, $|A| < 1\%$, $|O| < 1\%$. This is a basic assumption in machine learning. In contrast, the original protocol has: $|M| = 76\%$, $|F| = 24\%$, $|B| = 50\%$, $|W| = 50\%$, $|O| = |A| = |H| = 0\%$ for training, and $|M| = 87\%$, $|F| = 13\%$, $|B| = 84\%$, $|W| = 12\%$, $|H| = 4\%$, $|A| < 1\%$, and $|O| < 1\%$ for testing. (b) The number of training images available is 4 times larger (about 45,000 vs 10,530 images), which might improve generalization. (c) Since realistic applications might involve age estimations from unknown persons, we restrict the test images to persons not appearing in the training data. In [42], 44% of the test images belong to persons seen in training. (d) One can improve evaluation accuracy by repeating the protocol several times (here 5 times). The original protocol is repeated twice.

The alternative protocol has more training images. Thus, we distort each image only 4 times to create the DR-dataset and once to create the S-dataset. Notice that all races appear in the training and test data. Therefore, during training we consider two classes, B and a virtual class R=W+A+H+O to preserve a binary clustered graph for race estimation and balance the size of the classes. However, for better comparison, to compute the CR for race only the B and W races are considered. The results are shown in Table 9.

7. Discussion

In this article, we propose an extension of HGSFA called hierarchical information-preserving GSFA (HiGSFA) that complements the slowness principle with information preservation improving global slowness, input reconstruction and estimation accuracy for supervised learning problems.

We analyze the advantages and limitations of HSFA (and HGSFA) networks, particularly the phenomena of unnecessary information loss and poor input reconstruction. Unnecessary information loss occurs when a node in the network prematurely discards information that would have been useful for slowness maximization in another node higher up in the hierarchy. Poor input reconstruction

Algorithm	Age (MAE)	Age (RMSE)	Race (CR)	Gender (CR)
CNN (5CV)*‡ [50]	3.88	—	—	—
BIF+3SVM †‡ [51]	4.2	—	—	—
HiGSFA (G_3)	3.511 ± 0.002	4.626 ± 0.026	98.80% ± 0.01	98.53 % ± 0.02
HiGSFA (age only)	$\mathbf{3.502 \pm 0.003}$	$\mathbf{4.583 \pm 0.000}$	—	—
Chance level	9.16	10.78	87.53%	80.69%

Table 9: Accuracy of HiGSFA on MORPH-II using the alternative protocol. Some results are written in this table as $a \pm b$, where a is the average over 5 runs (test data) and b is the standard error of the mean. *5-fold cross validation. †Used a larger version of the database but only 10 001 training images. ‡As in the alternative protocol, it was ensured that the persons in the training and test data are disjoint.

refers to the difficulty of approximating an input accurately from its feature representation. We show that these phenomena are the result of locally optimizing slowness, yielding suboptimal global features.

To address these shortcomings, we improve feature extraction at the local level. The feature vectors computed by the iGSFA nodes of HiGSFA are divided in two parts, a slow and a reconstructive part. The features of the slow part follow a slowness optimization goal and are slow features transformed by a linear scaling. The features of the reconstructive part follow the principle of information preservation (i.e. maximization of mutual information between outputs and labels), which we implement in practice as the minimization of a reconstruction error. A parameter $\Delta_T \approx 2.0$ balances the lengths of the slow and reconstructive parts, J and $D - J$ features respectively, where D is the output dimensionality and J is the number of features with $\Delta < \Delta_T$.

We also present a new method to combine various efficient training graphs into a single efficient training graph. With the combined graph we can learn several (numerical or categorical) labels simultaneously. Its construction is simple; one only has to add the node and edge weights of the graphs, and the only requirements are that the graphs are consistent and have the same or proportional node weights.

The experimental results show that HiGSFA is better than HGSFA in terms of feature slowness, input reconstruction and age estimation accuracy. Moreover, HiGSFA offers higher accuracies than current state-of-the-art algorithms for age estimation, including approaches based on BIF features and convolutional neural networks. The improvement is a reduction by 48.5 days ($\approx$ 1.5 months) in the average estimation error. This is a relatively small but significant improvement.

In the next sections, we discuss a few conceptual aspects of the proposed approach, the results and future work.

7.1. The approach

Information preservation can be guaranteed by preserving the information contained in the data that describes the global slow features. However, we show

that one cannot always identify this information at a local level. Therefore, we resort to the reconstruction goal and preserve as much information of the local input as possible, which is likely to also include information relevant to extract the global slow parameters.

We addressed feature garbling only briefly and are aware that this is a complex problem that needs better formalization. HiGSFA may only partially reduce this problem by the use of reconstructive features, which might be simpler than slow features. However, feature garbling is still present in the slow part.

The features extracted by HiGSFA are better than those of HGSFA quantitatively and qualitatively. Even if unlimited training data and computational resources were available, the features extracted by HGSFA do not necessarily converge to those of HiGSFA. In this overfitting-free scenario, information loss would only decrease partially in HGSFA, because the main cause of this problem is not overfitting but the local optimization of slowness.

There is a subtle similarity between the overall unsupervised learning of HiGSFA and some techniques for training CNNs (unsupervised pre-training and joint optimization of supervised and unsupervised objectives). These approaches result in more information from the input reaching the top nodes.

Another way to combine the slowness principle with information preservation is to optimize a single objective function that integrates both criteria, favoring directions that are slow and have a large variance. However, we found in previous experiments that balancing the two criteria is difficult in practice.

HiGSFA inherits from SFA a close connection to unsupervised learning. Like GSFA can be emulated with SFA (see [10]) HiGSFA can be approximated by training HiSFA unsupervisedly with data generated from a particular Markov chain. This emulation would incur in a small error due to PCA being unaware of sample weights, which could be fixed by using the node weights as weighting factors of the samples during the computation of the covariance matrix by PCA.

7.2. Network parameters

By selecting the network structure appropriately, the computational complexity of HiGSFA (and other hierarchical versions of SFA) is linear w.r.t. the number of samples and their dimensionality, offering feasible training times. Training a single HiGSFA network (231 660 images of 96×96 pixels) takes only 10 hours, whereas HiSFA takes about 6 hours (including the time needed for data loading, the supervised step, and the evaluation) on a single computer (24 cores and 128 GB of RAM) without GPU computing. However, the algorithm can be implemented on GPUs or using distributed processing, because the nodes in a layer may be trained independently. For comparison, the system of [42] takes 24.5 hours (training and testing).

HiGSFA is more accurate than HGSFA with the same output dimensionalities and network hyper-parameters. However, HiGSFA yields even higher accuracies if larger output dimensionalities are used than those of the HGSFA network. This can be explained by various factors: (a) In both networks, the input dimensionality of SFA is I' (the expanded dimension), whereas in HiGSFA

the input dimensionality of PCA is I, where typically $I' \gg I$. Hence, slow features may overfit more than reconstructive features for this setup. (b) In HGSFA, the features of all the nodes are attracted to the same optimal free responses, whereas in HiGSFA, the reconstructive part is attracted to the (different for each node) local principal components. Thus, in HGSFA overfitting might accumulate through the layers more than in HiGSFA. (c) The HGSFA network may benefit less from more input features, because faster features might be noise-like and result in more overfitting without providing much additional information.

In the iGSFA algorithm, the number of slow features J with $\Delta < \Delta_T$ should be smaller than the output dimensionality D, so that $D - J$ output features are reconstructive. Otherwise, the output features would not have a reconstructive part, removing the advantages of the method. This might happen if too many slow parameters, their mixtures, or higher-frequency harmonics are present in the data. One might avoid this problem by setting a smaller Δ_T, by directly controlling the number of the slow features, and by using training graphs with a small number of optimal free responses with $\Delta < 2.0$.

7.3. Age, gender and race estimation

We choose the problem of age estimation from adult facial photographs, because it appears to be an ideal problem to test the capabilities of HiGSFA. For age estimation, PCA is not very useful because wrinkles, skin texture and other higher-frequency features are poorly represented. Therefore, it is not obvious and even counter-intuitive that feature slowness improves by incorporating PCs in HiGSFA. Improvements on feature slowness using other supervised learning problems, such as gender, race or horizontal-position estimation, would be inconclusive because for such problems a few PCs code the discriminative information relatively well.

To estimate age, race and gender simultaneously, we propose a graph G_3 that combines three graphs, encoding sensitivity to the particular labels and favoring invariance to any other factor. The graphs used are a serial graph for age estimation with 32 groups and two clustered graphs with 2 clusters each for race and gender estimation. The number of features with $\Delta < 2.0$ that can be extracted from these graphs is 15, 1 and 1, respectively (due to their particular geometry). Since the node weights of these graphs are not proportional exactly, the combined graph has more than 17 optimal responses with $\Delta < 2.0$. Still, the first 5 features contain most of the relevant information, and the method works well in practice.

The reconstructive part is not used in the supervised step, only the first 5 slow features of the output. This shows that HiGSFA and HGSFA concentrate the label information in the first features. One can actually replace the iGSFA node on the top of the HiGSFA network by a regular GSFA node, so that all features are slow, without affecting the performance. The superiority in age estimation of HiGSFA over HGSFA is thus not due to the use of principal components in the final supervised step but to the higher quality of the slow features.

The performance of HiGSFA for age estimation is the highest reported on the MORPH-II database with an MAE of 3.497 years. Previous state-of-the-art results are an MAE of 3.63 years using a multi-scale CNN [43] and 3.92 using BIF+rKCCA+SVM [42].

Even though our system performs slightly better than state-of-the-art algorithms, this was not our focus here. Our goal was to improve on HGSFA. Thus, our claim is that HiGSFA is better than HGSFA regarding feature slowness, input reconstruction and estimation accuracy.

7.4. Reconstruction from slow features

The experiments confirm that PCA is more accurate than HiGSFA at reconstruction using 75 features, which was expected because PCA features are optimal when reconstruction is linear. In HiGSFA 67 features are reconstructive (8 less than in PCA) and they are computed hierarchically (locally), in contrast to the PCA features, which are global. Thus, it is encouraging that the gap between PCA and HiGSFA at input reconstruction is moderate. In turn, HiGSFA is more accurate than HGSFA, because reconstruction is the secondary goal of HiGSFA, whereas HGSFA does not pursue reconstruction. The improved reconstruction capability of HiGSFA might facilitate certain applications, such as morphing.

Since the HiGSFA network implements a nonlinear transformation, nonlinear reconstruction algorithms are also reasonable. Nonlinear reconstruction might provide more accurate reconstructions in theory, but we have not been able to train such type of algorithms well enough to perform better on test data than the simpler global linear reconstruction algorithm. One algorithm for nonlinear reconstruction minimizes the feature error. However, since the number of dimensions is reduced by the network, one can expect many samples with feature error $e_{\text{feat}} = 0$ (or e_{feat} minimal) that differ substantially from any valid input sample, not looking at all like a face. To correct this problem, one might need to consider the input distribution to select an appropriate reconstruction. Although generative adversarial networks [52, 53, 54] were originally designed to generate random inputs, it might be possible to adapt them to do reconstruction from HiGSFA networks.

7.5. Future work

Estimation accuracy may be improved by using more complex hierarchical networks, for instance, by increasing the overlap of the receptive fields and using more complex nonlinearities. As usual, more training images might also improve accuracy, which could be approximated by implementing true face-distortion methods and not just simple transformations at the image level.

One key factor for the performance of multi-scale CNN is the use of receptive fields centered at specific facial points (compare with non-aligned receptive fields). This idea could also be applied to HiGSFA, and might particularly boost generalization.

We have informally tested HiGSFA on other problems (e.g., face detection, digit and traffic sign recognition) with good results. Therefore, we are interested in a formal and systematic evaluation of HiGSFA on these and more problems.

The idea of complementing GSFA with information preservation in hierarchical networks can also be applied to SFA (e.g., either using iSFA nodes or an HiGSFA network and a linear graph with samples ordered by time). The improved feature slowness of HiSFA over HSFA might be useful to improve simulations based on SFA for Neuroscience. Moreover, existing neural models based on the slowness principle might benefit from incorporating information preservation.

7.6. Final words

We believe it is possible to develop successful learning algorithms based on a few simple but strong learning principles and heuristics, and this is the approach that we try to pursue with HiGSFA. We are not so much interested in developing an algorithm that might be strong but cannot be understood analytically.

The slow and reconstructive parts of the extracted features can be seen as two information channels, the first one codes information connected to the slow parameters, and the second one codes information representing the input. Although the information in the slow part is somewhat mixed, it can be further decomposed into three channels. The 3rd slowest feature is mostly related to race, the 4th one to gender, and the remaining ones to age. Therefore, HiGSFA follows two of the suggestions in [55] based on findings from Neuroscience on the primate visual system for successful computer vision, namely, hierarchical processing and information-channel separation.

The proposed algorithm is general purpose (e.g., it does not know anything about face geometry), but it is still capable of outperforming special-purpose state-of-the-art algorithms, at least for the age estimation problem. This shows the improved versatility and robustness of the algorithm, and makes it a good candidate for many other problems of computer vision on high-dimensional data, particularly those lying at the intersection of image analysis, nonlinear feature extraction, and supervised learning.

References

[1] I. Rish, G. Grabarnik, G. Cecchi, F. Pereira, G. J. Gordon, Closed-form supervised dimensionality reduction with generalized linear models, in: Proc. of the 25th ICML, ACM, New York, NY, USA, 2008, pp. 832–839.

[2] G. E. Hinton, Connectionist learning procedures, Artificial Intelligence 40 (1-3) (1989) 185–234.

[3] P. Földiák, Learning invariance from transformation sequences, Neural Computation 3 (2) (1991) 194–200.

[4] G. Mitchison, Removing time variation with the anti-hebbian differential synapse, Neural Computation 3 (3) (1991) 312–320.

[5] L. Wiskott, Learning invariance manifolds., in: Proc. of 5th Joint Symposium on Neural Computation, San Diego, CA, USA, Vol. 8, Univ. of California, 1998, pp. 196–203.

[6] L. Wiskott, T. Sejnowski, Slow feature analysis: Unsupervised learning of invariances., Neural Computation 14 (4) (2002) 715–770.

[7] H. Sprekeler, C. Michaelis, L. Wiskott, Slowness: An objective for spike-timingdependent plasticity?, PLoS Comput Biol 3 (6) (2007) 1136–1148.

[8] V. R. Kompella, M. D. Luciw, J. Schmidhuber, Incremental slow feature analysis: Adaptive low-complexity slow feature updating from high-dimensional input streams, Neural Computation 24 (11) (2012) 2994–3024.

[9] A. N. Escalante-B., L. Wiskott, Slow feature analysis: Perspectives for technical applications of a versatile learning algorithm, Künstliche Intelligenz [Artificial Intelligence] 26 (4) (2012) 341–348.

[10] A. N. Escalante-B., L. Wiskott, How to solve classification and regression problems on high-dimensional data with a supervised extension of slow feature analysis, Journal of Machine Learning Research 14 (2013) 3683–3719.

[11] A. N. Escalante-B., L. Wiskott, Gender and age estimation from synthetic face images with hierarchical slow feature analysis., in: International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, 2010, pp. 240–249.

[12] A. N. Escalante-B., L. Wiskott, Theoretical analysis of the optimal free responses of graph-based sfa for the design of training graphs, e-print arXiv:1509.08329 (2015).

[13] N. M. Mohamed, H. Mahdi, A simple evaluation of face detection algorithms using unpublished static images, in: 10th International Conference on Intelligent Systems Design and Applications, 2010, pp. 1–5.

[14] R. A. Fisher, The use of multiple measurements in taxonomic problems, Annals of Eugenics 7 (1936) 179–188.

[15] M. Sugiyama, Local fisher discriminant analysis for supervised dimensionality reduction, in: Proc. of the 23rd ICML, 2006, pp. 905–912.

[16] W. Tang, S. Zhong, Computational Methods of Feature Selection, Chapman and Hall/CRC, 2007, Ch. Pairwise Constraints-Guided Dimensionality Reduction.

[17] D. Zhang, Z.-H. Zhou, S. Chen, Semi-supervised dimensionality reduction, in: Proc. of the 7th SIAM International Conference on Data Mining, 2007.

[18] M. Sugiyama, T. Idé, S. Nakajima, J. Sese, Semi-supervised local fisher discriminant analysis for dimensionality reduction, Machine Learning 78 (1-2) (2010) 35–61.

[19] H. Sprekeler, T. Zito, L. Wiskott, An extension of slow feature analysis for nonlinear blind source separation., Cognitive Sciences EPrint Archive (CogPrints):7056 (2010).

[20] H. Sprekeler, On the relation of slow feature analysis and laplacian eigenmaps, Neural Computation 23 (12) (2011) 3287–3302.

[21] M. Franzius, N. Wilbert, L. Wiskott, Invariant object recognition with slow feature analysis, in: Proc. of the 18th ICANN, Vol. 5163 of LNCS, Springer, 2008, pp. 961–970.

[22] P. Berkes, Pattern recognition with slow feature analysis, Cognitive Sciences EPrint Archive (CogPrints):4104 (2005).

[23] P. Koch, W. Konen, K. Hein, Gesture recognition on few training data using slow feature analysis and parametric bootstrap, in: International Joint Conference on Neural Networks, 2010, pp. 1 –8.

[24] T. Kuhnl, F. Kummert, J. Fritsch, Monocular road segmentation using slow feature analysis, in: Intelligent Vehicles Symposium, IEEE, 2011, pp. 800–806.

[25] Z. Zhang, D. Tao, Slow feature analysis for human action recognition, IEEE Transactions on Pattern Analysis and Machine Intelligence 34 (3) (2012) 436 –450.

[26] T. Zhang, D. Tao, X. Li, J. Yang, Patch alignment for dimensionality reduction, IEEE Transactions on Knowledge and Data Engineering 21 (9) (2009) 1299–1313.

[27] M. Franzius, H. Sprekeler, L. Wiskott, Slowness and sparseness lead to place, head-direction, and spatial-view cells, PLoS Computational Biology 3 (8) (2007) 1605–1622.

[28] M. Franzius, N. Wilbert, L. Wiskott, Invariant object recognition and pose estimation with slow feature analysis., Neural Computation 23 (9) (2011) 2289–2323.

[29] A. N. Escalante-B., L. Wiskott, Heuristic evaluation of expansions for Non-Linear Hierarchical Slow Feature Analysis., in: Proc. The 10th International Conference on Machine Learning and Applications, Los Alamitos, CA, USA, 2011, pp. 133–138.

[30] N. Wilbert, Hierarchical slow feature analysis on visual stimuli and top-down reconstruction, Ph.D. thesis, Humboldt-Universität zu Berlin, Mathematisch-Naturwissenschaftliche Fakultät I (2012).

[31] T. Blaschke, T. Zito, L. Wiskott, Independent Slow Feature Analysis and Nonlinear Blind Source Separation., Neural Computation 19 (4) (2007) 994–1021.

[32] T. Xia, D. Tao, T. Mei, Y. Zhang, Multiview spectral embedding, Trans. Sys. Man Cyber. Part B 40 (6) (2010) 1438–1446.

[33] T. Zito, N. Wilbert, L. Wiskott, P. Berkes, Modular toolkit for data processing (mdp): a python data processing framework, Frontiers in Neuroinformatics 2 (8).

[34] N. Ramanathan, R. Chellappa, S. Biswas, Age progression in human faces: A survey, Journal of Visual Languages and Computing 15 (2009) 3349 – 3361.

[35] Y. Fu, G. Guo, T. S. Huang, Age synthesis and estimation via faces: A survey, IEEE Transactions on Pattern Analysis and Machine Intelligence 32 (11) (2010) 1955–1976.

[36] X. Geng, Z.-H. Zhou, K. Smith-Miles, Automatic age estimation based on facial aging patterns, IEEE Transactions on Pattern Analysis and Machine Intelligence 29 (12) (2007) 2234–2240.

[37] G. Guo, G. Mu, Y. Fu, T. S. Huang, Human age estimation using bio-inspired features, in: CVPR, 2009, pp. 112–119.

[38] G. Guo, G. Mu, Y. Fu, C. Dyer, T. Huang, A study on automatic age estimation using a large database, in: IEEE 12th International Conference on Computer Vision, 2009, pp. 1986–1991.

[39] K. Luu, K. Ricanek, T. Bui, C. Suen, Age estimation using active appearance models and support vector machine regression, in: IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems, 2009, pp. 1–5.

[40] G. Guo, G. Mu, Human age estimation: What is the influence across race and gender?, in: IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 2010, pp. 71–78.

[41] G. Guo, G. Mu, Simultaneous dimensionality reduction and human age estimation via kernel partial least squares regression, in: CVPR, 2011, pp. 657–664.

[42] G. Guo, G. Mu, A framework for joint estimation of age, gender and ethnicity on a large database, Image and Vision Computing 32 (10) (2014) 761–770, best of Automatic Face and Gesture Recognition 2013.

[43] D. Yi, Z. Lei, S. Li, Age estimation by multi-scale convolutional network, in: Computer Vision – ACCV 2014, Vol. 9005 of Lecture Notes in Computer Science, Springer International Publishing, 2015, pp. 144–158.

[44] K. Ricanek Jr., T. Tesafaye, Morph: A longitudinal image database of normal adult age-progression, in: Proceedings of the 7th International Conference on Automatic Face and Gesture Recognition, FGR '06, IEEE Computer Society, Washington, DC, USA, 2006, pp. 341–345.

[45] S. Klampfl, W. Maass, Replacing supervised classification learning by Slow Feature Analysis in spiking neural networks, in: Proc. of NIPS 2009: Advances in Neural Information Processing Systems, Vol. 22, MIT Press, 2010, pp. 988–996.

[46] P. Berkes, Handwritten digit recognition with nonlinear fisher discriminant analysis, in: ICANN, Vol. 3697 of LNCS, Springer Berlin/Heidelberg, 2005, pp. 285–287.

[47] The FG-NET aging database.
URL `http://sting.cycollege.ac.cy/~alanitis/fgaging/index.html`

[48] B. Ni, Z. Song, S. Yan, Web image and video mining towards universal and robust age estimator, Multimedia, IEEE Transactions on 13 (6) (2011) 1217–1229.

[49] S. E. Choi, Y. J. Lee, S. J. Lee, K. R. Park, J. Kim, Age estimation using a hierarchical classifier based on global and local facial features, Pattern Recogn. 44 (6) (2011) 1262–1281.

[50] I. Huerta, C. Fernndez, C. Segura, J. Hernando, A. Prati, A deep analysis on age estimation, Pattern Recognition Letters.

[51] H. Han, C. Otto, A. Jain, Age estimation from face images: Human vs. machine performance, in: International Conference on Biometrics (ICB), 2013, pp. 1–8.

[52] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, in: Advances in Neural Information Processing Systems 27, Curran Associates, Inc., 2014, pp. 2672–2680.

[53] S. C. Alec Radford, Luke Metz, Unsupervised representation learning with deep convolutional generative adversarial networks, e-print arXiv:1511.06434 (November 2015).

[54] E. L. Denton, S. Chintala, a. szlam, R. Fergus, Deep generative image models using a laplacian pyramid of adversarial networks, in: Advances in Neural Information Processing Systems 28, Curran Associates, Inc., 2015, pp. 1486–1494.

[55] N. Krüger, P. Janssen, S. Kalkan, M. Lappe, A. Leonardis, J. Piater, A. Rodriguez-Sanchez, L. Wiskott, Deep hierarchies in the primate visual cortex: What can we learn for computer vision?, IEEE Transactions on Pattern Analysis and Machine Intelligence 35 (8) (2013) 1847–1871.